6

A Systems-Level Perspective on Attention and Cognitive Control

Guided Activation, Adaptive Gating, Conflict Monitoring, and Exploitation versus Exploration

Jonathan D. Cohen, Gary Aston-Jones, and Mark S. Gilzenrat

THE SCOPE OF ATTENTION

An understanding of attention is arguably one of the most important goals of the cognitive sciences and yet also has proven to be one of the most elusive. Most attention researchers will agree that a major problem has been agreeing on a definition of the term and the scope of the phenomena to which it applies. There are no doubt as many explanations for this state of affairs as there are those who consider themselves "attention researchers." However, most will probably agree that, in large measure, this is because attention is not a unitary phenomenon—at least not in the sense that it reflects the operation of a single mechanism, or a single function of one or a set of mechanisms. Rather, attention is the emergent property of the cognitive system that allows it to successfully process some sources of information to the exclusion of others, in the service of achieving some goals to the exclusion of others. This begs an important question: If attention is so varied a phenomenon, how can we make progress in understanding it? There are two simple answers to this question: Be precise about the specific (aspects of the) phenomena to be studied, and be precise about the mechanisms thought to explain them.

In this chapter, we address a particular type of attentional phenomenon—that associated with cognitive control. Furthermore, we focus on an account that addresses not only the functional characteristics of this form of attention but also how it is implemented in neural machinery. This neurally oriented approach is attractive not only because it is intrinsically interesting to understand how the mechanisms of the brain give rise to the processes of the mind but more specifically because this exercise has proven useful in generating insights into how controlled attention operates at the systems level. By assuming that information is

represented as patterns of activity, and information processing occurs as the flow of activity, it becomes possible to understand how information represented in one part of the system can influence the processing of information in other parts of the system—that is, how attention and control operate at the systems level. The sections that follow develop this idea in greater detail, first by providing a particular example of controlled attention and how it can be modeled in terms of explicit processing mechanisms, then by showing how it can explain some of the most important observations that have been made about attention and control, and finally by reviewing recent elaborations of the basic model that have begun to address broader questions about the psychological and neural mechanisms that underlie cognitive control.

AN EXAMPLE

To see how attentional effects can be understood in terms of specific processing mechanisms, it is useful to consider an example of a model of a specific task. There are now models of a variety of tasks that could serve this purpose well. Here, we focus on a model of the Stroop task (Cohen, Dunbar, & McClelland, 1990), because this task has occupied such a central role in studies of attention (both basic and clinical), because the model illustrates in a relatively straightforward manner the principles of interest, and because this model has been used to explain a wide array of findings using the Stroop task. This model was developed within the connectionist or parallel distributed processing (PDP) framework, which has been described and elaborated in great detail elsewhere (e.g., McClelland, 1993; O'Reilly & Munakata, 2000; Rumelhart, McClelland, & PDP Research Group, 1986), and therefore we assume it is familiar (or accessible) to the reader.

In the Stroop task (Stroop, 1935), subjects must attend to one dimension of a stimulus (e.g., the color in which a word is displayed) and ignore a competing but prepotent dimension (e.g., the word itself). For example, subjects are asked to name the color of an incongruent stimulus, such as the word *green* displayed in red. In our model, units are arranged into two pathways (Figure 6.1). Stimulus units representing the color project to associative units in the color-naming pathway, which project in turn to verbal response units. The word pathway converges on the same set of verbal response units. Furthermore, connections are stronger in the word pathway, capturing the assumption that written words are more frequently and consistently mapped to their pronunciations than are visual color stimuli to the utterance of their names. As a result, with no additions to the model, it will respond to the incongruent Stroop stimulus above by "reading" the word (i.e., activating the "green" response unit). In fact, this is how human subjects respond if not instructed otherwise. That is, they produce the strongest (e.g., most familiar or salient) response to a stimulus. Critically, however, they can respond to the weaker dimension of a stimulus when asked to do so (i.e., name the color in the Stroop task). This an elementary—and perhaps the most studied—form of controlled, or voluntary, attention.

To explain this ability, we make the following set of modifications to the model. First, we assume that at-rest units have relatively low activity. This corresponds well with the properties of neurons (especially those in cortical areas), which typically exhibit relatively low firing rates at rest. This can be seen in Figure 6.2a, by noting that for an input of zero, the activity of the unit is also near zero.[1] Second, we include an additional set of task demand units,[2] each of which corresponds to one of the tasks subjects are asked to perform (color naming and word reading). We assume that each of these units is connected to all of

the associative units in the corresponding pathway. Thus, the color-naming unit is connected to the associative units in the color-naming pathway, and similarly for the word-reading unit. When one of these task demand units is activated, it sends activity to all the associative units in the corresponding pathway. This has the effect of sensitizing these units to input from the stimulus units. Because this effect sits at the core of how attention operates in this model, it is important to consider this in greater detail.

Figure 6.2a illustrates the activation function for a unit—that is, the function that determines its activity based on the summed input it receives from other units. Note that this function is nonlinear.[3] This is central to our account of attentional effects. Recall that units have low activity at rest (that is, when their input is zero). In this range the activation function is relatively flat. In other words, even if one of the stimulus units were to be activated and pass activity to the corresponding associative unit, this would have limited impact on the activity of that associative unit. Now assume that one of the task-demand units is activated. This passes activity to the associative units in that pathway. Let us assume further that the amount of activity is sufficient to move these units to the midpoint of their activation function, where this is steepest. Note that this does not provide any specific information to that pathway. That is, all the units in that pathway have been equally activated, so none drive one response more than the others. However, now any input to these units from the stimulus units will have a large impact on their activity. Even a small excitatory input to one of these associative units will quickly drive its activity up, while inhibitory input to the other will drive its activity down. In other words, the effect of activating the task-demand unit is

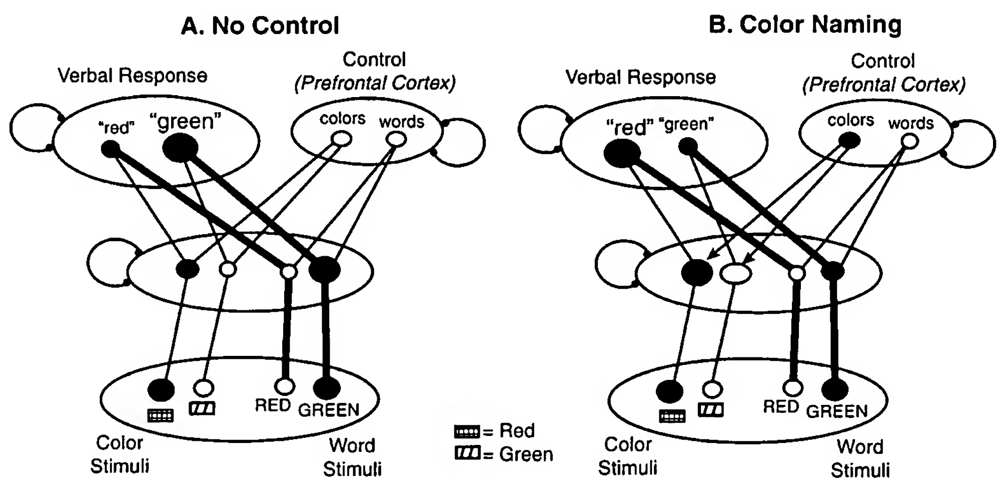

FIGURE 6.1. Model of the Stroop task. Circles represent processing units and line weights the strength of connections between units. Active units are filled (larger = more active). Looped connections with small black circles indicate mutual inhibition among units within that layer. (A) No control. Activation of conflicting inputs in the two pathways (green word input and red color input) produces a response associated with the word, due to the stronger connections in the word reading pathway. (B) Control present. The color task-demand unit is activated, representing the current intent to name the color. This passes activation to the associative units in the color naming pathway, which primes those units (indicated by larger size), and biases processing in favor of activity flowing along this pathway. This biasing effect favors activation of the response unit corresponding to the color input, even though the connection weights in this pathway are weaker than those in the word pathway.

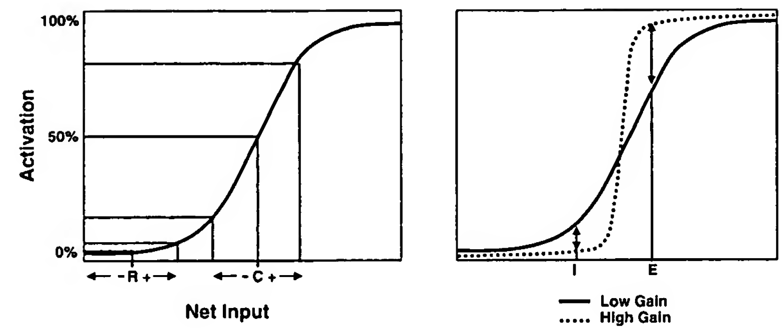

FIGURE 6.2. Activation function (see note 1 for the equation for this function). (A) Units are inhibited at rest (point labeled "R"), so that a change in the net input has relatively little effect on the unit's activity. Top-down input (from the task-demand, or control layer) places the unit near the midpoint of its activation function (point labeled "C"), where a change in the net input (+ = excitatory; – = inhibitory) has a considerably larger impact on the unit's activity. (B) An increase in gain (dotted line) increases the activity of units receiving excitatory input (E), and decreases the activity of units receiving inhibitory input (I), thus increasing the contrast between activated and inhibited units.

to bias the associative units in that pathway, placing them in the sensitive range of their activation function. This serves to modulate the responsivity of those units, making them more sensitive to the inputs.[4] This, in turn, allows the system to respond selectively to one source of information while ignoring another. For example, by activating the color-naming task-demand unit, the model can now respond to the color of the stimulus even when a conflicting word stimulus is present. That is, the model exhibits attention. This attentional effect derives from the ability of the task-demand units to guide the flow of activity along one pathway, while attenuating the flow along another. For this reason, we have come to refer to this as the guided activation theory of cognitive control (Miller & Cohen, 2001).

ATTENTION AND RELATED CONSTRUCTS

Models that implement the guided activation theory provide a quantitative account of attentional effects in a wide range of tasks (e.g., Braver, Barch, & Cohen, 1999; Braver & Cohen, 2001; Cohen, Servan-Schreiber, & McClelland, 1992; Cohen, Romero, Serven-Schreiber, & Farah, 1994; Dehaene & Changeux, 1989, 1991; Mozer, 1988; O'Reilly, Noelle, Braver, & Cohen, 2002; O'Reilly & Munakata, 2000; Phaff, van der Heijden, & Hudson, 1990; Servan-Schreiber, Bruno, Carter, & Cohen, 1998; Schneider & Detweiler, 1988). Equally important, it provides a unifying account of a constellation of processes and constructs related to attention. These are considered in the remainder of this section.

Controlled versus Automatic Processing

This is one of the oldest, and most fundamental constructs in cognitive psychology (Posner & Snyder, 1975; Shiffrin & Schneider, 1977). This distinction is cast largely in terms of the reliance on attention. Controlled processes are defined as those that rely on attention for execution, while automatic processes are defined as those that can be carried out without at-

tention. One of the earliest applications of this construct was to the Stroop task (Posner & Snyder, 1975). Color naming was considered to be controlled because it relies on attention. Without attention to the color, subjects will read the word. Furthermore, the color has no impact on word reading, even when it conflicts with the word being read. Conversely, word reading is automatic because it does not appear to rely on attention. Even when asked to name the color, if a conflicting word is present it slows the response to the color (the classic Stroop effect; MacLeod, 1991; Stroop, 1935). This is thought to reflect the fact that the word is processed even without the allocation of attention. However, there are problems with a simple dichotomous distinction between color naming and word reading in terms of controlled versus automatic processing.

First, it is not clear that any cognitive process can occur entirely independently of attention. For example, although an individual is reading the words on this page, presumably he or she is not doing so out loud. Thus word reading, at least as it is practiced in the Stroop task, is not entirely independent of attention and control. Second, it is not clear that color naming is always dependent on attention and control. In a clever experiment, MacLeod and Dunbar (1988) had participants learn associations between arbitrary shapes (displayed in black and white) and names for them that happened to be color words. At various points during training, they tested their participants' ability to name shapes that were displayed in colors that conflicted with their names. As might be guessed, they found that early in training a shape's color interfered with the ability to provide its name. In other words, in this task, color naming behaved as if it were the automatic process, contradicting the traditional suggestion that it is controlled. Kahneman and Treisman (1984) reviewed a number of other attentional findings, concluding that all processes rely on attention to some degree, and that this may vary in a graded fashion. This is consistent with MacLeod and Dunbar's findings (which also demonstrated that, as subjects became more proficient at shape naming, the color of the shape came to influence this less, while the shape's name came to interfere with color naming). Our model offers a mechanistically explicit account of these findings.

As noted earlier, in the absence of any input from the task-demand units, neither color naming nor word reading can be carried out. This is because the associative units in both pathways rest in too unresponsive a region of their activation function. Thus, even word reading requires attention. At the same time, connections in the word-reading pathway are stronger. Thus, the amount of activity of the task-demand unit needed to support word reading is less than that needed for color naming. In other words, word reading relies less on attention or control than does color naming. It is critical to note, however, that even in the absence of task-demand unit activity *some* information can flow along a pathway. Although this may not be enough to elicit an overt response, it may be enough to influence processing. Thus color naming can be influenced by information in the word pathway and therefore demands more attention when the word conflicts with the color name than when it is congruent. However, this effect is relative. For example, when color naming competes with a *weaker* process, the reverse will be true. This was the situation in MacLeod and Dunbar's (1998) shape-naming experiment, in which the association of a shape with a color word is weaker than a color with a color word. Thus, different processes vary in the degree to which they rely on attention, and this also varies for a given process based on the context in which it is carried out. From this perspective, the distinction between controlled and automatic processing is not dichotomous and absolute but, rather, graded and relative: Some processes are more automatic than others, and processes vary in their automaticity based on the context in which they occur. While this demands quantitative rather than qualitative characteri-

zation of processes, our model offers a framework in which such quantification can be carried out in terms of the connection strengths in the relevant pathways.

Types of Attention

We began the previous section by defining controlled processes as those that rely on attention. But how do we define attention? In our model, attention can be defined as the influence that activity in the task-demand units has on processing in the color and word pathways. Note that this does not rely on any qualitatively distinct mechanisms. Attention arises from the flow of activity between units and over connections that are qualitatively identical to those used to actually perform the task. This suggests that we can define attention, in the most general way, as the modulatory influence that representations of one type have on selecting which (or to what degree) representations of other types are processed, that is, how representations of one type guide the flow of activity among other types. Thus, the representation of an object may influence which sensory features are processed, just as in our model the representation of a task demand influenced which dimensions of the stimulus were processed. The representation of a strong stimulus may even have a "bottom-up" influence on representations of task demands.[5] This idea accords well with a theory of attention that has emerged from the neurophysiological literature—the biased competition theory proposed by Desimone and Duncan (1995). This theory assumes that in the brain, different representations compete for expression, and that the role of attention is to bias this competition in favor of some competitors over others. The source of the bias can be bottom up (e.g., driven by a stimulus) or top down (driven by a higher level "template"). In our model, the top-down flow of activity from the task-demand units literally biased the associative units in each pathway, modulating their responsivity and thereby influencing the competition between information in each pathway and guiding its flow along one relative to the other.

As appealing as the generality of this perspective on attention is, it raises an important issue. On the one hand, the guided activation theory and the biased competition theory emphasize the general nature of attention and the broad range of circumstances in which attentional interactions can occur. This provides a general framework for thinking about attention and, in our case, explicitly modeling attentional interactions. On the other hand, the very breadth of this range invites the following, as yet unanswered question: Are there meaningful distinctions to be made between attentional interactions that occur in different domains, at different levels of processing, or in different "directions" (e.g., top down vs. bottom up). For example, are there systematic differences in the dynamics or scope of modulatory interactions? One answer to this question was implicit in the introduction to this chapter, in the choice to focus on "controlled attention," suggesting that attentional interactions related to "control" exhibit a cohesive set of properties that distinguish them from other types of attentional interactions. This choice is motivated both by functional and neurobiological considerations and bears a close similarity to a highly influential taxonomy proposed by Posner and Petersen (1990; see also Posner, 1980). These authors distinguished between three attentional systems in the brain: an anterior attentional system (housed in the frontal lobes) associated with cognitive control and action selection, a posterior attentional system (housed in the parietal and occiptal lobes) associated with orienting and perceptual attention, and an arousal system (subserved by brainstem neuromodulatory systems) associated with sustained attention and vigilance. Work building on our model of attention and its relationship to neural mechanisms has reached a largely convergent perspective that makes

complementary contributions to our understanding of how attention and cognitive control operate and are implemented in the brain.

Functional Requirements for Cognitive Control

The modeling efforts previously discussed have focused on a specific set of mechanisms that explain how cognitive control gives rise to attentional effects. However, a more general consideration of cognitive control suggests that additional mechanisms are required for its operation. To see this, let us first consider the function of the task-demand units in the Stroop model. So far, we have focused on their attentional effects—that is, their ability to select one source of information for processing over another. However, more generally, they can be viewed as implementing a mapping from a particular set of inputs to a particular set of outputs. For example, the color task-demand unit represents the relationship between color stimuli and their names. From this vantage point, the task-demand units can be seen as carrying out the function of rules, intentions, or goal representations in other theoretical frameworks. All these specify a relationship between existing states (determined by external sensory inputs or internal influences such as memories, emotions, etc.) and desired outcomes that demand particular behaviors in order to be achieved.

Two critical requirements for the representation of task demands, rules, intentions, or goals are that these be actively maintained while the relevant behaviors are performed and then adaptively updated when behavior has achieved the desired outcome or is no longer appropriate: the task is complete, the rule has changed, or the intention or goal has been achieved. This suggests that the apparatus responsible for cognitive control include mechanisms responsible for active maintenance and adaptive updating. Interestingly, these functions are just those for which the prefrontal cortex (PFC) appears to be specialized. This has led us to propose that the PFC subserves the function carried out by the task-demand units in our model: The active and sustained representation of task demands, such as rules and goals—what we have sometimes referred to more broadly as internal representations of context (Cohen, Braver, & O'Reilly, 1996; Cohen & Servan-Schrieber, 1992)—that bias, or guide, the flow of activity along task-relevant pathways, in accord with the guided activation theory. However, the original model lacks critical features, such as the ability to determine on its own which task-demand representation should be active, just how active this should be, for how long, and how this should be updated when a new one is required—that is, it lacks mechanisms for adaptive updating.

Work over the past decade has directly addressed these issues by augmenting the basic model, constrained by neuroscientific data. Although a detailed consideration of these developments is beyond the scope of this chapter, a brief review illustrates how cognitive control can be implemented in a neural system that is self-organized and self-regulated, without recourse to unexplained mechanisms or intelligence (i.e., without the need for a "homunculus").

A NEURAL SYSTEM FOR CONTROLLED ATTENTION

Active Maintenance

The first requirement for a system of control is that it be able to actively maintain representations of task demand, rules, or goals over temporally extended periods (e.g., during performance of the task), in the absence of external support. For example, subjects do not need

to be reminded, trial after trial, to name the color of the Stroop stimulus in a block of trials. Models of sustained activity in PFC have implemented this property using recurrent connectivity (Braver, Cohen, & Servan-Schreiber, 1995; Durstewizt, Seamans, & Sejnowski, 2000; Wang, 1999; Zipser, 1991). This gives rise to attractor dynamics, allowing sets of units with mutually excitatory connections to actively maintain themselves in the absence of input (i.e., as an "attractor state"; Hopfield & Tank, 1986). Recently, biophysically more detailed explorations have suggested that intracellular mechanisms may also contribute to active maintenance, allowing individual units to be "latched" into an on or an off state (Frank, Louhry, & O'Reilly, 2001). However, the importance that this has for models at the systems level remains to be explored.

Adaptive Updating

A second critical property is that the system must be able not only to maintain task representations but also to update these appropriately. This requires that representations in PFC resist perturbation by task-irrelevant inputs (i.e., avoid distraction), while responding to inputs that signal the need for a change (i.e., avoid perseveration). The ability for appropriate updating is central to the flexibility of cognitive control, and disturbances of PFC are known to be associated with distractibility, perseveration, or both. There is growing evidence to support the hypothesis that updating relies on a dopamine-mediated adaptive gating mechanism subserved by the ventral tegmental area (VTA), a dopaminergic nucleus in the brainstem that projects widely to prefrontal areas. In initial work (Braver & Cohen, 2000; O'Reilly, Braver, & Cohen, 1999), we evaluated the plausibility of a simple version of this hypothesis, by implementing a transient gating signal in the task-demand (PFC) layer, that rendered these units temporarily responsive to input from posterior structures (see Figure 6.3). In the absence of this gating signal, representations in the task-demand layer were insensitive to exogenous input, allowing them to maintain the current task demand representation against impinging sources of interference. However, when a gating signal occurred, inputs from other parts of the system could drive activity in the task-demand layer, activating a new representation in that layer. More recently, we have begun to explore the possibility that a more powerful system involving dopaminergic projections to the basal ganglia subserves this gating function (e.g., Frank et al., 2001). This elaboration allows more focused forms of gating that can support hierarchical updating of goals, subgoals, and so on. Nevertheless, the proposal that a dopaminergic, brainstem-mediated gating mechanism regulates the updating of goal representations in PFC raises a critical question: How does this system know when to produce a gating signal, and what new state to produce in the PFC?

The answer to this question comes from important new discoveries regarding the effects of dopamine. Once thought to mediate the hedonic value of a reward, recent work suggests that dopamine release may function as a learning signal, reinforcing associations that provide better predictions of reward (Shultz, Dayan, & Montague, 1997). Importantly, the parameter used to implement this function in computational models (Montague, Dayan, & Sejnowski, 1996) bears a remarkable similarity to the parameter we have used in models to implement a dopamine-based gating signal (Braver & Cohen, 2000). In a series of models, we have illustrated that implementing concurrent effects of the dopamine signal on reinforcement learning and gating allows the system to associate stimuli with the gating signal that predict reward, and thus learn how to update representations in the PFC appropriately (e.g., Braver & Cohen, 2000; Rougier & O'Reilly, 2002). We have used these mechanisms to account for detailed behavioral and neurobiological data regarding the function of the

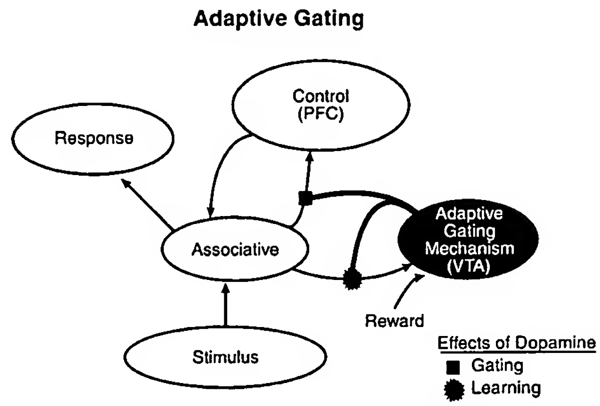

FIGURE 6.3. Model with an adaptive gating mechanism (VTA). Activity of the VTA regulates input to the PFC layer from the associative layer, at the same time training connections from the associative layer back to the VTA. Training occurs according to the temporal differences (predictive Hebbian) learning algorithm (Sutton, 1988). This compares inputs from the associative layer with reward signals, and strengthens connections from the associative layer to the VTA for cues that successfully predict reward. This interacts synergistically with the gating mechanism, because cues that are associated with PFC representations that lead to the procurement of reward thereby predict reward. This strengthens their connections to the VTA, so that in the future they will be more likely to generate a gating signal, activate their associated PFC representation, better predict reward, and therefore receive greater strengthening, and so on.

PFC (O'Reilly et al., 2002; Reynolds & Braver, 2002) in tasks that rely on the flexible deployment of control.

Conflict Monitoring and Regulation of the Degree of Control

The mechanisms described thus far address the ability to represent task-demand information, maintain it, and update it as needed. However, they do not address a different set of questions, which is how the system knows when top-down control is needed in the first place, and just how much is needed to achieve a goal? The need for such mechanisms becomes apparent if we assume that, as in the rest of the brain, representations within the PFC compete for expression. Such competition might provide a partial explanation for the capacity limitations of cognitive control. Such capacity limitations are readily apparent to anyone who has tried to attend to e-mail and talk on the phone at the same time or to talk to a passenger while driving under adverse conditions. These limitations are also fundamental to the distinction between controlled and automatic processing, which assumes that controlled processing relies on a capacity-limited attentional system (e.g., the use of dual-task designs to identify reliance on controlled processing rests on this assumption). Given competition within the PFC, it becomes important for the system to determine whether PFC representation is needed to perform a given task, and if so, how active the representation must be to support adequate performance. Insofar as a weakly active PFC representation leaves room for others to also be active (and thereby other goals to be simultaneously pursued), it is advantageous for the system to titrate the activity of PFC representations to current task needs. In our work with the original Stroop model, we examined the requirements that different processes had for top-down control. As discussed earlier, weaker pathways (e.g., for color

naming) required stronger activation of the corresponding task-demand unit to achieve the same level of performance, and this was especially true when conflicting information was present in competing pathways. Conversely, stronger pathways needed less (but always some) top-down support. Recently we have begun to consider what mechanisms might support the adaptive regulation of task demand activity to meet current task needs.

Our approach to this problem was inspired by the observations about attention with which we began: The primary function of attention is to support the processing of task-appropriate sources of information against competition from interfering sources. Put another way, the role of attention is to reduce *conflicts* in processing. Therefore, the occurrence of conflict provides a natural signal of the need for attentional control We have proposed that monitoring for conflicts in processing is subserved by a specific neural system: the anterior cingulate cortex (ACC). This hypothesis is more fully elaborated in Botvinick, Braver, Yeung, Ullsperger, Carter, and Cohen (Chapter 7, this volume). In brief, the conflict monitoring hypothesis argues that the ACC is responsive to conflict in processing pathways—in particular those that are subject to attentional control by the PFC (i.e., mappings from inputs to outputs)—and that activity in the ACC signals the need to more strongly activate representations in the PFC, in order to better support processing in the task-relevant pathway(s). For example, if the color-naming unit is not sufficiently activated in the Stroop model, then information flowing along this pathway will not compete effectively with interfering information arriving from the word pathway (in the case of an incongruent stimulus). Both response alternatives will become activated, and conflict will ensue. The conflict-monitoring hypothesis asserts that such conflict will engage the ACC, signaling the need to increase activity of the color-naming unit (see Figure 6.4). This hypothesis, and predictions that derive from it, have now received considerable support from behavioral and neuroscientific findings (for a review, see Botvinick et al., Chapter 7, this volume).[6]

The conflict-monitoring hypothesis makes the strong claim that while the ACC plays a role in conflict monitoring and the recruitment of attentional control, it is not responsible for the *allocation* of control. This function is ascribed to the PFC. This assertion contrasts with specific claims about the role of the ACC made by Posner and Petersen (1990) in their original formulation of the anterior attentional system (see also Posner & Dehaene, 1994; Posner & DiGirolamo, 1998). However, it does not violate the more general spirit of their proposal that top-down attentional control is subserved by a frontal system involving the ACC. The primary thrust of the conflict-monitoring hypothesis is to add further specification to one component of this system and suggest a modified set of structure–function relationships. One limitation of the conflict-monitoring hypothesis, however, is that in its present form it does not precisely characterize the mechanisms by which ACC conflict monitoring engages PFC control. The models that have been developed to date implement this as a direct influence of the ACC on the PFC. More recently, we have begun to explore alternative mechanisms that may mediate this influence. One hypothesis is that this involves the locus coeruleus.

The locus coeruleus (LC) is the brainstem neuromodulatory nucleus responsible for most of the norepinephrine (NE) released in the brain (Berridge & Waterhouse, 2003; Foote, Bloom, & Aston-Jones, 1983). It has widespread projections throughout the neocortex. Previous work (Servan-Schreiber, Printz, & Cohen, 1990) has suggested that the effects of NE release can be modeled as a change in the gain (steepness) of the activation function of connectionist units (see Figure 6.2b). This has the effect of augmenting the activity of

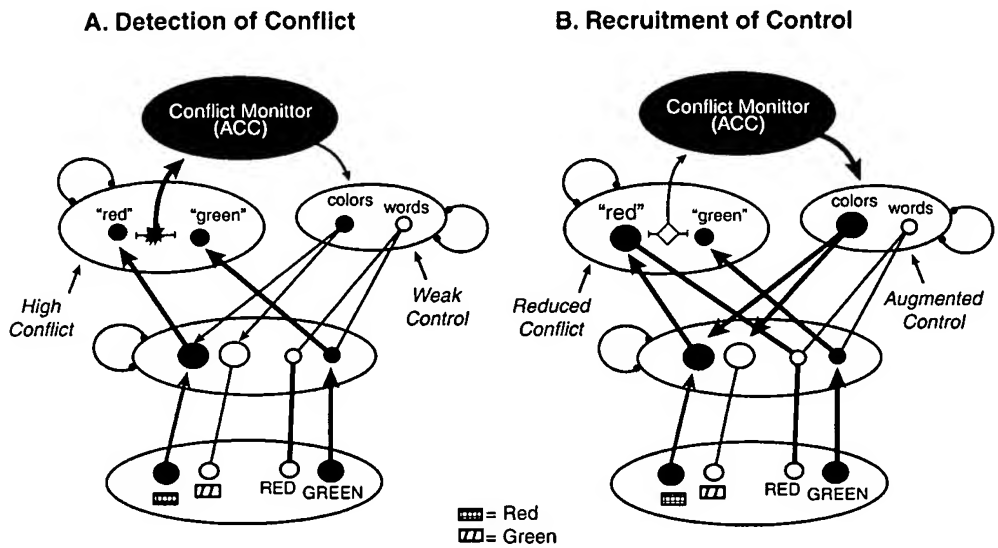

FIGURE 6.4. Model with conflict-monitoring mechanism (ACC). (A) Weak control allows conflict to develop at the response layer as a result of competing inputs from different pathways. This is detected by the ACC. (B) Detection of conflict within the ACC augments activated units in the PFC layer, producing additional top-down control and a reduction of conflict.

units that are already activated and further suppressing the activity units that are already being inhibited. That is, an increase in gain increases the contrast of the current pattern of activity. This produces precisely the effect within the PFC that is needed in response to conflict. In fact, such contrast enhancement of PFC representations is how adjustments in control have been implemented in our simulations of the conflict-monitoring hypothesis to date (Botvinick, Braver, Carter, Barch, & Cohen, 2001; Yeung, Botvinick, & Cohen, in press). However, in these models, contrast enhancement in the PFC was assumed to be produced directly by ACC activity. While this remains a possibility, we have begun to consider the possibility that this effect is actually mediated by NE release from the LC.

LC-mediated modulation of the PFC is consistent with several lines of evidence, including modeling work that specifies a role for the LC in attentional modulation (e.g., Gilzenrat et al., 2002; Holmes, Nieuwenhuis, & Gilzenrat, & Cohen, 2002; Robertson, Mattingley, Rorden, & Driver, 1998; Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999; Yu & Dayan, 2002), as well as recent neuroanatomic evidence suggesting that the ACC is a primary source of cortical projections to the LC (Rajkowski, Lu, Zhu, Cohen, & Aston-Jones, 2000). The appeal of this hypothesis is that it provides a mechanism by which conflict detection within the ACC can augment control without specific knowledge about which particular representations in PFC require augmentation: A global signal can have specific effects. Although details concerning the dynamics of LC-mediated modulation of the PFC are beyond what can be considered here, it is important to note that our hypothesis is that transient (phasic) rather than sustained (tonic) release of NE mediates the modulation of PFC. This forms part of a more general theory about the role of the LC in regulating attention, that we briefly review next.

Exploration versus Exploitation and Regulation of the Focus of Attention and Control

The control mechanisms described previously explain how behavior can be guided by representations of task demands, rules, intentions, or goals in the PFC, adaptively modulating their activity as needed to support task performance. According to the conflict-monitoring hypothesis, when performance degrades and conflict increases, control should increase. However, what about situations in which behavior continues to fall short, despite compensatory adjustments in control? For example, in a Stroop experiment, what if the color stimulus is progressively degraded? The gradual increase in conflict should lead to concomitant increases in control. However, at some point, if the color is degraded beyond recognizability, it makes no sense to further augment control. Rather, control should be withdrawn from this task and some other goal should be pursued. Or, consider the following situation that is perhaps more ecologically valid: An animal is picking berries from a tree. At first berries are everywhere, and the task may not require much effort or attention, but as berries become more scarce, more attention is needed. After some point, however, increasing attention will not help; there are just too few berries left to make the effort worthwhile. At this point, some other behavior should be pursued. These situations suggest that as conflict increases and reward diminishes, at some point the relationship between conflict and control should reverse. This tension between optimizing control to reap the benefits of the current behavioral program and abandoning the current program when it may be more advantageous to sample alternative behavioral programs is well recognized in machine learning, where it is referred to as the trade-off between exploitation and exploration.

We have hypothesized that interactions between the ACC and the LC may regulate the balance in this trade-off (Usher et al., 1999). This builds on observations made by Aston-Jones, Rajkowski, Kubiak, and Alexinsky (1994), indicating that in the awake behaving monkey, the LC shifts between two operating modes that correspond closely with behavioral performance in a simple target detection task. In the "phasic mode," when the animal is performing optimally, the LC shows only moderate levels of tonic discharge, while it exhibits a phasic response selectively to target stimuli but not to distractors. In the "tonic mode," the tonic level of discharge is higher, but there are no phasic responses to target stimuli, reaction time to targets is slower, and the animal commits a greater number of false alarms to distractors.[7] Usher and colleagues (1999) developed a biophysically plausible model of LC function that accounted for transitions between phasic and tonic modes in terms of a single physiological variable (coupling between LC cells) and at the same time explained the impact of these shifts on task performance.[8] In brief, the model suggests that in the phasic mode, the strong LC response to target stimuli facilitates their processing by transiently increasing gain which, at the response layer, has the effect of lowering response threshold. Because this effect is selective to targets, it improves both target detection performance and reaction time. In contrast, in the tonic mode, the absence of a phasic response accounts for the increase in response time to target stimuli, while the increased but indiscriminant release of NE lowers the system's response threshold to other stimuli (e.g., distractors), accounting for the increase in false alarms.

The specificity of the LC phasic response to the target (in the phasic mode) has been demonstrated in a number of additional studies, including reversal experiments in which LC has been shown to reliably acquire the new target before the animal's overt behavior has done so (e.g., Aston-Jones, Rajkowski, & Kubiak, 1997). These characteristics suggest that

in the phasic mode, the LC functions as an attentional filter that selects for the occurrence (i.e., *timing*) of task-relevant stimuli, much as cortical attentional systems filter the *content* of a stimulus. The existence of such a temporal filter is consistent with several recent psychophysical studies (e.g., Coull & Nobre, 1998). At the same time, by increasing the gain of cortical representations, the LC phasic response can also enhance the effects of cortical selection by content (e.g., the top-down effects of the PFC). Together, these effects allow the LC phasic response to selectively facilitate responses to task-relevant stimuli when they occur. One immediate question, however, is, What is the adaptive value of the tonic mode, which seems to deteriorate performance? We have argued that the tonic mode may support exploratory behavior. By reducing phasic responses to target stimuli, and increasing tonic NE release, the system is more effectively driven by task-irrelevant stimuli. Such responsiveness is not adaptive with respect to the specified experimental task, because it permits the processing of task-irrelevant stimuli and the sampling of other behavioral programs. However, it may be highly adaptive if either the current task is no longer remunerative, or if the environment changes and more valuable opportunities for reward or new behavioral imperatives have appeared. From this perspective, a shift from the LC phasic to the tonic mode may shift the behavioral strategy from exploitation (when this is no longer adaptive) to exploration (when a new goal should be sought). Loosely speaking, this "throws the ball up in the air, so another team can take it." Viewed from the perspective of attentional control, the LC phasic mode supports the current control state (exploitation), while the LC tonic mode provokes a withdrawal of control from the current task, favoring the sampling of other behavioral goals (exploration), which raises one more important question: What information can the system use to determine whether it should exploit (LC phasic mode) or explore (LC tonic mode)? The answer to this question closes the loop, both in the control system and in our discussion.

We hypothesize that ACC conflict monitoring can provide the necessary information. This hypothesis requires only one additional assumption: that the ACC is able to integrate conflict over two time frames—a short one (seconds) and a long one (minutes). Consider the following two circumstances. In one, performance is good and there are still rewards to be accrued from the current task, but there are occasional lapses in performance producing transient increases in conflict (e.g., on single trials). Under these conditions, control should be increased each time there is a lapse, to restore performance. That is, control should be increased when long-term conflict is low but there is a momentary increase in short-term conflict. In contrast, consider a different situation in which performance has been poor and conflict has been *persistently* high. At some point, this situation should encourage the withdrawal of control, irrespective of short-term changes in conflict. A similar situation should arise when, irrespective of performance, opportunities for reward diminish. A relatively simple equation can capture these relationships,[9] which indexes the need for control as a function of short-term conflict and reward, discounted by the accumulation of long-term conflict and dimunition of reward (see Figure 6.5). This computation can be used to drive shifts between LC phasic and tonic modes by influencing simple physiological parameters (Brown et al., 2004; Usher et al., 1999). Taken together, these mechanisms would constitute a self-regulating system that is responsive to demands for control over different time scales and is sensitive to the current value of exploration versus exploitation (Figure 6.6). These mechanisms are consistent with known properties of the LC, ACC, and their anatomic connectivity; however, their validation presents a challenge to further neurophysiological investigation.[10]

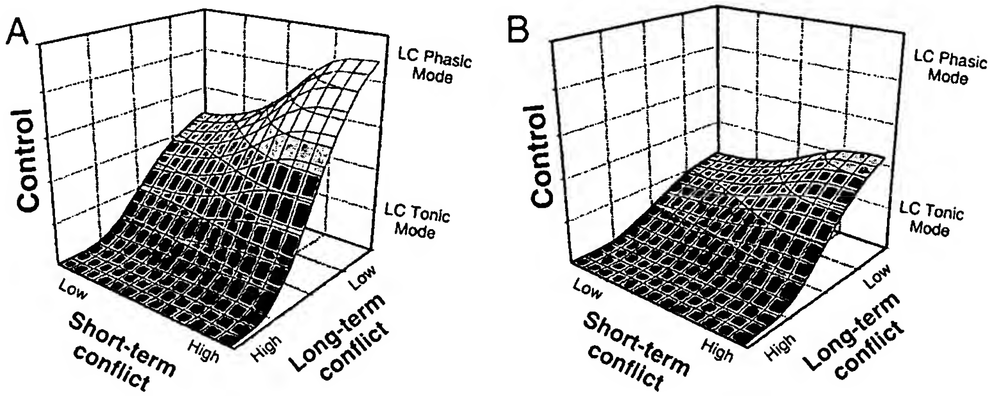

FIGURE 6.5. Relationship of control to conflict and reward (see note 9 for the equation defining this relationship). (A) Under conditions of high reward and low long-term conflict, transient increases in conflict (short-term conflict) produce greater control, by driving LC into the phasic mode. This effect is diminished as long-term conflict increases, driving LC into tonic mode. (B) Low reward damps the effect of short-term conflict, driving the LC into the tonic mode irrespective of conflict.

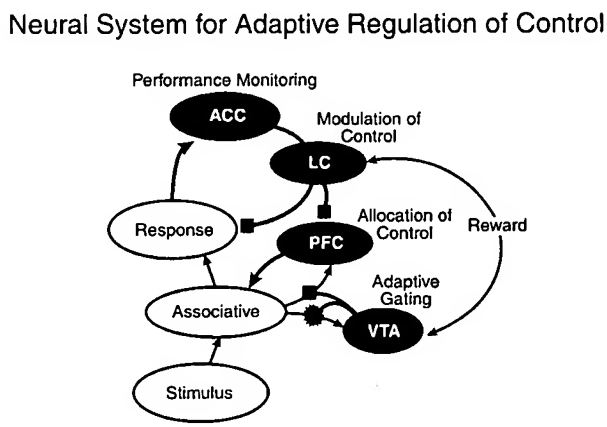

FIGURE 6.6. Integrated neural system for the adaptive regulation of control. This system is able to adaptively update control representations in PFC, learn how to do so, and modulate the strength of top-down control in response to prevailing balance between reward and short-term and long-term degradations in performance (as indexed by conflict).

CONCLUSIONS AND FUTURE CHALLENGES

This chapter has reviewed recent theoretical work, pursued within the connectionist framework, that addresses the nature of controlled attention and its neural implementation. The goal of this work is to provide a mechanistically explicit account of how processes that give rise to the phenomena of attention and control are implemented in the brain. The hypotheses reviewed suggest that control relies on the activation of appropriate representations in the PFC. These representations can be thought of as task demands, rules, intentions, or goals that direct behavior to produce desired outcomes by biasing processing and guiding the flow of activity along pathways responsible for mapping inputs to desired outputs. As summarized in Figure 6.6, PFC representations are regulated by several mechanisms, including a dopamine-mediated system for updating PFC representations in specific task contexts, and learning how to do so; an ACC-mediated system for assessing the demand for control; and an LC-mediated system for modulating PFC representations in response to these demands. These hypotheses define a mechanistically explicit, self-regulating system of control that is responsive to adaptive needs at different time scales and of fundamentally different types (e.g., exploration vs. exploitation). One important feature of these hypotheses is their suggestion that brainstem neuromodulatory systems—once thought to be responsible for the regulation of nonspecific aspects of psychological function, such as motivation (dopamine) and arousal (NE)—may play a significantly more central and specific role in information processing.

Of course, these mechanisms represent only a first step toward a more complete understanding of the neural mechanisms that underlie attention and control. First, they address only controlled attention, and not the many other forms and levels of attentional effects. However, even within the scope of controlled attention, many challenges remain. For example, it is possible to register a goal or intention for the future and then to dispatch this for pursuit sometime (hours, days or even years) in the future. These forms of control cannot be explained by active maintenance of representations in PFC alone, but they are likely to involve interactions between the PFC and medial temporal lobe structures that subserve episodic memory (Cohen & O'Reilly, 1996). There are also critical interactions between control and motivation, most likely involving interactions between the PFC and limbic structures. Perhaps the most perplexing puzzle that remains concerns the nature of representations in the PFC. Our models to date have stipulated the presence of representations in the PFC required to perform a given task. However, it seems unlikely that without infinite capacity, the PFC can house all the possible representations needed to meet the arbitrarily large set of potential task demands. A characteristic feature of human behavior is the flexibility of control, manifest as the ability to perform novel tasks, or to creatively structure new forms of behavior in a novel task environment. How then, with a large but finite set of resources, can the system exhibit the flexibility we witness in our everyday behavior? This question is closely related to an equally important one: How do representations develop in PFC?

Recent work has begun to address many of the questions raised here (for reviews, see Miller & Cohen, 2001; O'Reilly et al., 1999, 2002; Rougier, Noelle, Braver, Cohen, & O'Reilly, 2004). Nevertheless, the human ability to flexibly deploy attention and control, navigating the vast repositories of information available both from the environment and the system's stored knowledge, to respond appropriately under familiar circumstances and creatively under unfamiliar ones, remains one of the most fundamental and interesting myster-

ies of science. We hope to have illustrated in this chapter that this mystery need not remain intractable to theoretical analysis.

ACKNOWLEDGMENTS

The work reviewed in this chapter has been conducted under the generous and ongoing support of the National Institute of Mental Health and the National Alliance for Research on Schizophrenia and Depression. The authors would also like to note that most of the ideas and findings described in this chapter have emerged from longstanding collaborations with Deanna Barch, Matthew Botvinick, Todd Braver, Cameron Carter, Sander Nieuwenhuis, David Noelle, Randy O'Reilly, and Nick Yeung, as well as Eric Brown and Philip Holmes.

NOTES

1. This reflects the assumption that units are negatively biased. This can be seen from the expression for the activation function: activity = $1/(1 + e^{-(\text{net input})})$. For a net input of 0, this evaluates to an activity of 0.5. However, by assuming that a constant negative bias value is added to the net input (e.g., –4 as was used in Cohen et al., 1990), the unit will have a low activity value for a zero net input.
2. Task-demand units have, in different papers—both our own and those of others—variously been referred to as attention units, context units (designating the "task context"), mapping units, and goal units, which reflects the various perspectives that can be taken on their function (to which we return at the end of the chapter). It also highlights the perils of pursuing scientific research using natural language and the value of explicit, formal models. Ultimately, our theory of attention and its relationship to cognitive control is expressed in the form of the model and *not* in the words that are used to describe it. Thus, where any ambiguity or disagreement may arise about what is claimed by our theory, the model itself (and whatever elaboration of it is required to address the issue) is the final arbiter and not any words used to describe it.
3. The assumption of nonlinearity is justified on both theoretical and empirical grounds. Theoretically, it has been shown that using nonlinear rather than linear units confers considerably greater computational power on the system (e.g., Rumelhart, Hinton, & Williams, 1986). Empirically, it is clear that neurons (whether considered individually or in populations) have finite upper and lower bounds on their activity levels.
4. Two points are worth noting about this effect. First, the term "modulation" often implies a multiplicative effect. Here, however, modulation is produced by adding (biasing) activity from the task-demand unit to the associative units. Nevertheless, because this addition occurs to the net input, which appears in the exponent of the activation function, the effect is in fact multiplicative. Second, note that the effect of this modulation is to place units in the linear range of their activation function, which provides a justification for mathematical models that assume linear response properties (e.g., Brown & Holmes, 2001), by suggesting that those units primarily responsible for processing (i.e., those that are in the "focus of attention") reside in the linear range of an otherwise nonlinear response function.
5. This assumes that there are bottom-up connections to the task-demand units in the model. Although these were not included in the original Cohen et al. (1990) model, variants have included such connections (e.g., Cohen & Huston, 1994; O'Reilly & Munakata, 2000).
6. While this hypothesis proposes that conflict monitoring is subserved by the ACC, it does not claim that this is the *only* function of this structure. Rather, we view conflict monitoring as one of a family of functions subserved by the ACC, that monitor internal processing states for breakdowns in performance, much as the amygdala is thought to monitor the environment for external signs of threat.

7. At present, it is not clear whether these represent dichotomous modes, or ends of a continuum of states that the LC can occupy.
8. Recently, Brown and colleagues (2004) have proposed that changes in baseline firing rate may also serve to drive transitions between LC phasic and tonic modes. Which of these mechanisms (electrotonic coupling, baseline firing rate, or both) is actually operative in the LC remains an area of inquiry. However, what is relevant for present purposes is that LC mode can be determined by one or two easily regulated physiological parameters.
9. Control = reward / $[f(\text{conflict}_{\text{long-term}})*(1 - f(\text{conflict}_{\text{short-term}}))]$, where $f(\text{conflict}) = 1/(1 + e^{-\text{conflict}})$
10. The LC receives extensive projections from both the ACC and the orbitofrontal cortex (Aston-Jones et al., 2002; Rajkowski et al., 2000), which may provide evaluative information regarding both performance and rewards (e.g., Bush et al., 2002; Rolls, 2000). We should also note that the LC has extensive projections throughout the brain (except the hypothalamus and striatum, which it does not innervate). We propose that while NE release in PFC directly modulates control representations, simultaneous release of NE in areas outside the PFC serves to reinforce this effect in other sensory, motor, and associative areas.

REFERENCES

Aston-Jones, G., Rajkowski, J., & Kubiak, P. (1997). Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience, 80*, 697–715.

Aston-Jones, G., Rajkowski, J., Kubiak, P., & Alexinsky, T. (1994). Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *Journal of Neuroscience, 14*(7), 4467–4480.

Aston-Jones, G., Rajkowski, J., Lu, W., Zhu, Y., Cohen, J. D., & Morecraft, R.J. (2002). Prominent projections from the orbital prefrontal cortex to the locus coeruleus in monkey. *Society for Neuroscience Abstracts, 28*, 86–89.

Berridge, C. W., & Waterhouse, B. D. (2003). The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Research: Brain Research Review, 42*, 33–84.

Botvinick, M. M., Braver, T. S., Carter, C. S., Barch, D. M., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108*(3), 624–652.

Braver, T. S., Barch, D. M., & Cohen, J. D. (1999). Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. *Biological Psychiatry, 46*(3), 312–328.

Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 713–737). Cambridge, MA: MIT Press.

Braver, T. S., & Cohen, J. D. (2001). Working memory, cognitive control, and the prefrontal cortex: Computational and empirical studies. *Cognitive Processing, 2*, 25–55.

Braver, T. S., Cohen, J. D., & Servan-Schreiber, D. (1995). A computational model of prefrontal cortex function. In D. S. Touretzky, G. Tesauro, & T. K. Leen (Eds.), *Advances in neural information processing systems* (Vol. 7, pp. 141–148). Cambridge, MA: MIT Press.

Brown, E., & Holmes, P. (2001). Modeling a simple choice task: Stochastic dynamics of mutually inhibitory neural groups. *Stochastics and Dynamics, 1*(2), 159–191.

Brown, E., Moehlis, J., Holmes, P., Clayton, E., Rajkowski, J., & Aston-Jones, G. (2004). *The influence of spike rate and stimulus duration on noradrenergic neurons.* Manuscript submitted for publication.

Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., et al. (2002). Dorsal anterior cingulate cortex: a role in reward-based decision making. *Proceedings of the National Academy Science USA, 99*(1), 523–528.

Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: Recent developments and current challenges. *Philosophical Transactions of the Royal Society of London B, (Biological Sciences), 351, 1515–1527.*

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review, 97*(3), 332–361.

Cohen, J. D., & Huston, T. A. (1994). Progress in the use of parallel distributed processing models for understanding attention and performance. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 453–476). Cambridge, MA: MIT Press.

Cohen, J. D., & O'Reilly, R. C. (1996). A preliminary theory of the interactions between prefrontal cortex and hippocampus that contribute to planning and prospective memory. In M. Brandimonte, G. O. Einstein, & M. A. McDaniel (Eds.), *Prospective memory: Theory and applications* (pp. 267–295). Hillsdale, NJ: Erlbaum.

Cohen, J. D., Romero, R. D., Servan-Schreiber, D., & Farah, M. J. (1994). Mechanisms of spatial attention: The relation of macrostructure to microstructure in parietal neglect. *Journal of Cognitive Neuroscience, 6*(4), 377–387.

Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review, 99*, 45–77.

Cohen, J. D., Servan-Schreiber, D., & McClelland, J. L. (1992). A parallel distributed processing approach to automaticity. *American Journal of Psychology, 105*, 239–269.

Coull, J. T., & Nobre, A. C. (1998). Where and when to pay attention: The neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *Journal of Neuroscience, 18*, 7426–7435.

Dehaene, S., & Changeux, J. P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience, 1*(3), 244–261.

Dehaene, S., & Changeux, J. P. (1991). The Wisconsin card sorting test: Theoretical analysis and modeling in a neuronal network. *Cerebral Cortex, 1*(1), 62–79.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18*, 193–222.

Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *Journal of Neurophysiology, 83*(3), 1733–1750.

Foote, S. L., Bloom, F. E., & Aston-Jones, G. (1983). Nucleus locus ceruelus: New evidence of anatomical and physiological specificity. *Physiology Review, 63*, 844–914.

Frank, M. J., Loughry, B., & O'Reilly, C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience, 1*(2), 137–160.

Gilzenrat, M. S., Holmes, B. D., Holmes, P. J., Rajkowski, J., Aston-Jones, G., & Cohen, J. D. (2002). A modified Fitzhugh-Nagumo system simulates locus coeruleus-mediated regulation of cognitive performance. *Neural Networks, 15*, 647–663.

Holmes, B. D., Nieuwenhuis, S., Gilzenrat, M. S., & Cohen, J. D. (2002). *The role of locus coeruleus in mediating the attentional blink: A neurocomputational model.* Poster presented at the annual meeting of the Psychonomic Society, Kansas City, MO.

Hopfield, J. J., & Tank, D. W. (1986). Computing with neural circuits: A model. *Science, 233*, 625–633.

Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuranam & D. R. Davies (Eds.), *Varieties of attention* (pp. 29–59). New York: Academic Press.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109*(2), 163–203.

MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory and Cognition, 14*(1), 126–135.

McClelland, J. L. (1993). Toward a theory of information processing in graded, random, and interactive networks. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies*

in experimental psychology, artificial intelligence, and cognitive neuroscience (pp. 655–688). Cambridge, MA, MIT Press.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, *16*(5), 1936–1947.

Mozer, M. (1988). A connectionist model of selective attention in visual perception. In *Proceedings of the tenth annual conference of the Cognitive Science Society* (pp. 195–201). Hillsdale, NJ: Erlbaum.

O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically-based neural network model of working memory. In P. Shah & A. Miyake (Eds.), *Models of working memory* (pp. 375–411). Cambridge, UK: Cambridge University Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

O'Reilly, R. C., Noelle, D. C., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex in dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex*, *12*, 246–257.

Phaf, R. H., Van der Heijden, A. H. C., & Hudson, P. T. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology*, *22*(3), 273–341.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25.

Posner, M. I., & Dehaene, S. (1994). Attentional networks. *Trends in Neuroscience*, *17*, 75–79.

Posner, M. I., & DiGirolamo, G. J. (1998). Executive Attention: Conflict, target detection and cognitive control. In R. Parasuraman (Ed.), *The attentive brain* (pp. 401–423). Cambridge: MIT Press.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, *13*, 25–42.

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition* (pp. 55–85). Hillsdale, NJ: Erlbaum.

Rajkowski, J., Lu, W., Zhu, Y., Cohen, J., & Aston-Jones, G. (2000). Prominent projections from the anterior cingulate cortex to the locus coeruleus (LC) in rhesus monkey. *Society of Neuroscience Abstracts*, *26*, 838.15.

Reynolds, J. R., & Braver, T. S. (2002, April). *Computational and neural mechanisms of task-switching*. Paper presented at the ninth annual meeting of the Cognitive Neuroscience Society, San Francisco.

Robertson, I. H., Mattingley, J. B., Rorden, C., & Driver, J. (1998). Phasic alerting of neglect patients overcomes their spatial deficit in visual awareness. *Nature*, *395*, 169–172.

Rolls, E. T. (2000). The orbitofrontal cortex and reward. *Cerebral Cortex*, *10*(3), 284–294.

Rougier, N. P., Noelle, D. C., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (2004). *Prefrontal cortex and the flexibility of cognitive control: Rules without symbols*. Manuscript in submission.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 318–362). Cambridge, MA: MIT Press.

Rumelhart, D. E., McClelland, J. L., & PDP Research Group. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Schneider, W., & Detweiler, M. (1988). The role of practice in dual-task performance: toward workload modeling in a connectionist/control architecture. *Human Factors*, *30*(5), 539–566.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.

Servan-Schreiber, D., Bruno, R., Carter, C., & Cohen, J. D. (1998). Dopamine and the mechanisms of cognition. Part I: A neural network model predicting dopamine effects on selective attention. *Biological Psychiatry*, *43*, 713–722.

Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, *249*, 892–895.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning automaticity, attending and a general theory. *Psychological Review, 84*(2), 127–190.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643–662.

Sutton, R. (1988). Learning to predict by the method of temporal difference. *Machine Learning, 3*, 9.

Usher, M., Cohen, J. D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999). The role of locus coeruleus in the regulation of cognitive performance. *Science, 283*, 549–554.

Wang, X. J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *Journal of Neuroscience, 19*, 9587–9603.

Yeung, N., Botvinick, M. M., & Cohen, J. D. (in press). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychologial Review.*

Yu, A. J., & Dayan, P. (2002). Acetylcholie in cortical inference. *Neural Networks, 15*, 719–730.

Zipser, D. (1991). Recurrent network model of the neural mechanism of short-term active memory. *Neural Computation, 3*(2), 179–193.

Guilford Publications
72 Spring Street
New York, NY 10012
212-431-9800
800-365-7006
www.guilford.com

2171524

13 The Skill of Attention Control: Acquisition and Execution of Attention Strategies

Daniel Gopher

13.1 PROBLEM AND SCOPE

How able and efficient are humans in controlling and utilizing their limited mental resources? Can they be taught to improve their attention management skills? Can attention control be treated as a basic skill component? These are the main questions that will be examined in the present chapter. To support the notion of attention management as a trainable skill, we need evidence to show (1) the existence of behavioral potential, namely, control over the allocation of attention, (2) difficulties and failure in fulfilling this potential, and (3) ability to overcome or diminish these difficulties with proper training.

We seek data that bear upon the ability of performers to cope efficiently with tasks that require them to divide attention and processing efforts among multiple, dynamically varying elements. Such requirements are common in many daily tasks of intermediate and high complexity, where multiple dynamic elements have to be monitored, responded to, and interacted with to achieve a common general goal. For example, car drivers are required to divide attention among the manual control of the vehicle, monitoring the road, the traffic signs, and the behavior of other vehicles, while also attempting to locate themselves geographically in the environment. A similar task with much-elevated demands is that of a pilot who moves at high speed through six degrees of freedom space. Other examples include the task of a process controller in the control room of a modern power plant and a basketball player running with the ball while looking for the best opportunity to attempt a shot or pass, while also trying to anticipate defensive moves.

Common to all these tasks is that the performer would gain most if he or she could fully attend and respond to all elements at all times. However, such full attention is not possible. Hence, some trade-offs and priorities must be established along with attention-allocation strategies. Setting priorities is a common human experience; the question is how competent we are in establishing attention strategies and allocating processing efforts among concurrently changing task elements.

A *strategy* in this context is defined as a vector of differential weights or attention biases assigned to task elements. It influences the performer's mode

of response to the requirements of the task (see also Logan 1985). A strategy represents the solution developed by the performer to cope better with task demands and performance objectives, within the boundaries of his or her processing and response limitations.

For example, several subcomponents of driving were listed, all of which covary concurrently and should be attended to while driving. However, their relative importance changes considerably (i.e., a different strategy has to be adopted) when the driver tries to arrive earlier at his destination, searches for an address in an unfamiliar territory, or pays extra caution to driving regulations after noticing a police car parked at the side of the road. Moreover, there are several ways in which each of these objectives can be pursued, and each may require a different attention strategy. Also, strategies and the importance of elements may vary in the course of driving, as the situation or the intentions of the driver change.

Two major questions can be posed regarding the adoption and execution of attention strategies. One concerns how cognizant performers are about the efficiency of their investments, and about the benefits and costs of alternative strategies. A second question concerns how able humans are in the control and mobilization of their processing efforts. Related to these questions are issues such as the trainability of attention control and the best methods of training.

With few exceptions (e.g., Logan 1985), strategic control of attention along the lines described in the previous examples has never been a main and independent topic of research in experimental psychology. Although fuzzy notions of a central executive and attention-control mechanism have been incorporated in many models of the human processing system (Gopher and Sanders 1984; Kahneman 1973; Norman and Shallice 1981), little systematic research has been conducted to explore the nature and capabilities of such a controller. Manipulations and evidence bearing on the issue of control have been embedded in the pursuit of other interests, such as unitary limited capacity, processing resources, mental workload, and modes of processing. Nonetheless, the literature is quite rich with experimental paradigms and data demonstrating successful strategic control. There is also ample evidence for control failures and deficient knowledge about the efficiency of investments.

A major claim of the present chapter is that from both theoretical and practical vantage points, strategic control of attention is an important topic that merits independent status in experimental psychology. To support this claim I will first briefly review evidence portraying successful control and problems of attention control. This review is followed by a description of the main results from two lines of experiments conducted in our laboratory, where acquisition and transfer of attention management skills have been studied in the context of performing complex tasks. The final section draws some principles and conclusions, outlining directions for future research.

13.2 EVIDENCE ABOUT ATTENTION CONTROL

Demonstrations of Successful Control

Experimental data demonstrating successful voluntary control of attention include results from the tasks of focusing, dividing, and switching attention together with investing graded amounts of processing efforts.

Focusing Attention The most widely researched and documented aspect of voluntary attention control is the ability of humans to adopt a selective attention set and combat interference from irrelevant information. Research over the last four decades has shown that subjects can focus, lock on, and prioritize almost any feature that distinguishes one string of stimuli from another consistently. Although some features of the stimulus such as its spatial location and physical characteristics appear to be easier to focus upon and follow, voluntary selection is highly flexible and strategic. Throughout the years, this has led researchers to theoretical debates concerning early versus late selection models of attention limitations (Kahneman 1973; Kahneman and Treisman 1984). It is possible that evidence for one or more of these models may reflect the strategic freedom of the central executive in the selection of information.

For example, recent experiments on visual attention have provided elegant demonstrations of the separation between the "physical eye," dominated by the physical properties of the stimuli and the anatomy and physiology of the eye, versus the top-down selective-attention processes controlling the "mind's eye" (e.g., Posner, Snyder, and Davidson 1980). Contemporary research instigated by spotlight and zoom lens metaphors of visual attention has further demonstrated two strategic modes of focusing in the visual field, one general and the other specific mode (Castiliano and Umiltà 1990; Egeth 1977; Ericksen and St. James 1986; Jonides 1983). Under the general mode, attention is allocated evenly, in parallel, across the whole visual field. Under the specific mode, attention is focused on one display location, facilitating processing there, while inhibiting it at all others. Similar distinctions have been proposed by Kahneman (1973) for the auditory modality, when studying the behavior of subjects in dichotic listening tasks of selective attention.

Dividing Attention Data that demonstrate the successful division of attention come primarily from experiments with the dual-task paradigm (Gopher and Donchin 1986). Within this paradigm, subjects have to cope with the demands of concurrently presented tasks. In different versions of the paradigm, two time-shared tasks have to be treated as equally important, or one is designated as primary and the other secondary. A secondary task should be attended to only insofar as performance of the primary task is fully achieved.

The dual-task paradigm has been a major methodological tool in experiments on processing and response limitations, and on the properties of tasks that interfere or can be time-shared in concurrent performance (Gopher 1990;

Wickens 1980). The study of voluntary control of attention has been incidental to this research. Nonetheless, successful compliance with the equal-importance and secondary-task emphases are clear examples of attention control.

Switching Attention Switching attention from one task to another is an additional aspect of voluntary control that has been widely employed in the study of attention limitations. As before, the main interest in most studies has not been in the properties and control of switching, but rather in contrasting serial and parallel models of the human processing system. Isolated studies over the years attest to the significance of studying switching behavior and its control. One line of studies has documented consistent individual differences in attention-switching capabilities as a factor distinguishing between good and bad performers of complex tasks (Gopher 1982; Gopher and Kahneman 1971; Kahneman, Ben Ishai, and Lotan 1974; North and Gopher 1976).

Other support comes from the studies conducted by Posner and his colleagues (Posner and Rothbart 1986). They substantiate a distinction between two processes: focusing attention on a location in the visual field and processing information at that location, and disengaging and moving attention from one target to another. Experiments with brain-damaged subjects have led these researchers to identify different brain regions for the two processes. Most recently, Allport and Styles (1990) have conducted a series of experiments to test the ability of subjects to alternate and switch attention between tasks progressing serially through a single stimulus string. Their subjects demonstrated impressive ability to switch and adopt new processing and response sets in a wide variety of situations, while also exhibiting some difficulties that will be discussed in the next section.

Investing Graded Levels of Effort The evidence most related to a notion of attention strategies comes from experiments in which attention control was manipulated as part of empirical efforts to construct Performance Operating Characteristic (POC) functions. A POC is a curve depicting all possible levels of joint performance in two concurrent tasks, which arise by dividing into different priorities a limited resource that they must share. The theoretical analysis of POCs was proposed by Norman and Bobrow (1975) and elaborated by Navon and Gopher (1979) to distinguish between data and resource constraints on the human processing system. The main objective has been to characterize performance decrements under time-sharing conditions not only in terms of task parameters and structure, but also the allocation policy adopted by a performer and the availability of resources. Moreover, manipulation of strategic control has emerged as an important source of information about the nature and limitations of the central processor.

Empirical POCs have been constructed by manipulating the relative emphasis on two concurrent tasks through verbal instructions, payoff matrices, and on-line augmented feedback (e.g., Gopher, Brickner, and Navon 1982; Sperling and Melchner 1978; Spitz 1988; Strayer and Kramer 1990; Wickens and

Gopher 1977). Although, as before, the focus of this research was not the ability to control attention, but rather the mapping of central-processor limitations, it has produced monotonic POCs comprising three, five, and seven different intertask priority levels. This evidence supports the ability of subjects to adopt and successfully apply graded levels of investment.

Summary Although attention control has rarely been a direct topic of systematic research, several lines of experimental data support the ability of performers to adopt a selective set, divide and switch attention at will, and produce graded priority levels for two concurrently performed tasks.

Control Problems and Failures

Along with the positive evidence on voluntary attention control, there have been several reports that point to problems, failures, and limits of control.

Failure to Protect Primary-Task Performance A frequent finding in dual-task experiments, with one task as primary, is that subjects fail, despite clear instruction, to protect its performance. The introduction of a secondary task and responses to it are accompanied by performance decrements on the primary task. Such decrements have usually been interpreted to reflect capacity overload or structural interference between tasks (Gopher and Donchin 1986). However, they may also be manifestations of deficient control. Subjects may not have adequate knowledge about the attention costs of performing each of the tasks, or they may be unable to exercise sufficient attention control such that only "spare capacity" is allocated to the secondary task, leaving primary task performance intact. Also, these two types of control deficiencies may not be mutually exclusive.

Dissociation between Subjective Estimates and Performance Measures of Workload Research on the issue of measuring mental workload in demanding tasks has documented many cases of dissociation between the difficulty and demand estimates given by subjects, and the measures of load and demands derived from their actual task performance (Gopher and Braune 1984; Ye and Wickens 1988). The distribution of dissociations is quite symmetrical. Cases in which subjective estimates are higher than performance-based measures are as prevalent as the reverse order. Assuming that subjects made their best effort to perform the tasks, these dissociations may again imply either that subjects are not aware of the attention requirements of performance or that they invest their efforts inefficiently.

The Need for Augmented Feedback Although within the POC paradigm subjects have been shown to produce five and seven step monotonic POCs through investing graded levels of effort, these demonstrations have depended on the use of augmented feedback. Instructions for graded priority adjustments had to be accompanied by special, on-line, augmented feedback, displaying the

consequences of emphasis changes (e.g., Gopher, Brickner, and Navon 1982). Without such feedback or sufficient training, subjects were unable to perform multistep adjustments (Spitz 1986, 1988).

Attention Capture by Automatic Components In their study of controlled and automatic processes, Schneider and Fisk (1982) reported that under time-sharing conditions, task elements that were automated after prolonged training still captured attention, although controlled attention was not required to assure performance. Dedicated training efforts were needed to teach subjects to relax and release attention. This is another indication that either voluntary control or knowledge were deficient.

The Ability to Release Resources Voluntarily The problem hinted at by Schneider and Fisk (1982) has been revealed in a different form by our experiments with the POC paradigm. We have found that in many dual-task situations, subjects had as much difficulty lowering their standard of performance for the task for which priority was reduced as they had in improving performance for the task on which priority was increased. This phenomena was examined experimentally by Spitz (1986, 1988).

For example, when performing a tracking task concurrently with a letter-typing task, his subjects had difficulty complying with the requirement to lower their typing performance when the priority of this task was reduced, thereby allowing higher performance on the tracking task for which priority was increased. This phenomenon was equally powerful in difficult and easy tasks, precluding interpretations based on floor effects. Moreover, similar problems have been revealed in single-task conditions where emphasis instructions were given without any simultaneous competing task. Elsewhere we labeled this phenomena the problem of maintaining "minimal control levels" (Gopher 1982). It appears that subjects have difficulties lowering performance and reducing efforts on one task. They cannot easily release resources for the performance of another high-priority task, while still maintaining minimal control over the low-priority task.

Summary We have reviewed five types of problems and failures in the control of attention. In all types, the sources of difficulty appear to stem from insufficient knowledge about the efficiency of allocation or from deficient control capabilities. These deficits constitute clear targets for training.

Effects of Practice

Evidence concerning the contribution of practice to the improvement of attention control is scattered throughout the experimental literature. Most relevant studies have been conducted within the dual-task paradigm, using tasks whose performance posed a difficult or impossible mission for novice performers. With practice, time-sharing performance of the same tasks approximated the single-task levels of each. Typical examples are studies on the

ability of subjects to read while taking oral dictations (Hirst et al. 1980), on simultaneous piano playing and reading (Allport, Antonis, and Reynolds 1972), and on performing two independent search tasks (Schneider and Fisk 1984).

All these studies have been concerned primarily with the effects of practice on time-sharing performance. Improved attention control was not among the factors considered. Practice effects were primarily attributed to a general reduction of attention costs and to improved integration and organization of the tasks. Nonetheless, it is reasonable to assume, based on the evidence presented in the previous section, that control problems contributed to subjects' initial degraded performance, and improved control was part of the effects of training.

More support for the effects of practice on improved attention control comes from the experiments of Spitz (1986, 1988) on the POC paradigm. He showed that subjects can improve their minimal control levels through practice under both single- and dual-task conditions. That is, with practice, subjects are better able to comply with a requirement to lower their standard of performance and release resources for the performance of a concurrent task under time-sharing conditions. Other direct support, using a different experimental paradigm, comes from the work of Allport and Styles (1990). They reported effects of practice on improving the efficiency of attention switching to alternate continually between two or four tasks.

In conclusion, this review of recent studies substantiates the three prerequisites for a skill-oriented approach to study attention control. We have seen positive evidence for the existence of behavioral potential, several types of control difficulties, and indications (both direct and indirect) for improvement of control with practice. The following sections review two lines of studies conducted in our laboratory to investigate training of attention control and management.

13.3 DEVELOPING THE SKILL OF ATTENTION CONTROL

Methodology and Theoretical Approach

Our approach to the training of attention control relies on observations of subjects' time-sharing performance in the POC paradigm. Mainly, their difficulties lie in establishing multiple emphasis levels and in failing to adapt comfortably to the consequences of minimal control levels. We have found numerous direct and indirect indications of deficient knowledge and attention-control problems (e.g., Gopher and North 1977; Wickens and Gopher 1977). They have led us to reason that the capability of subjects to cope with concurrent task demands can be improved if they are taken through a well-designed sequence of relative priority changes. With such experience, they may acquire better representations for the value of differential investments and improve their ability to control them. These abilities may be applied in later

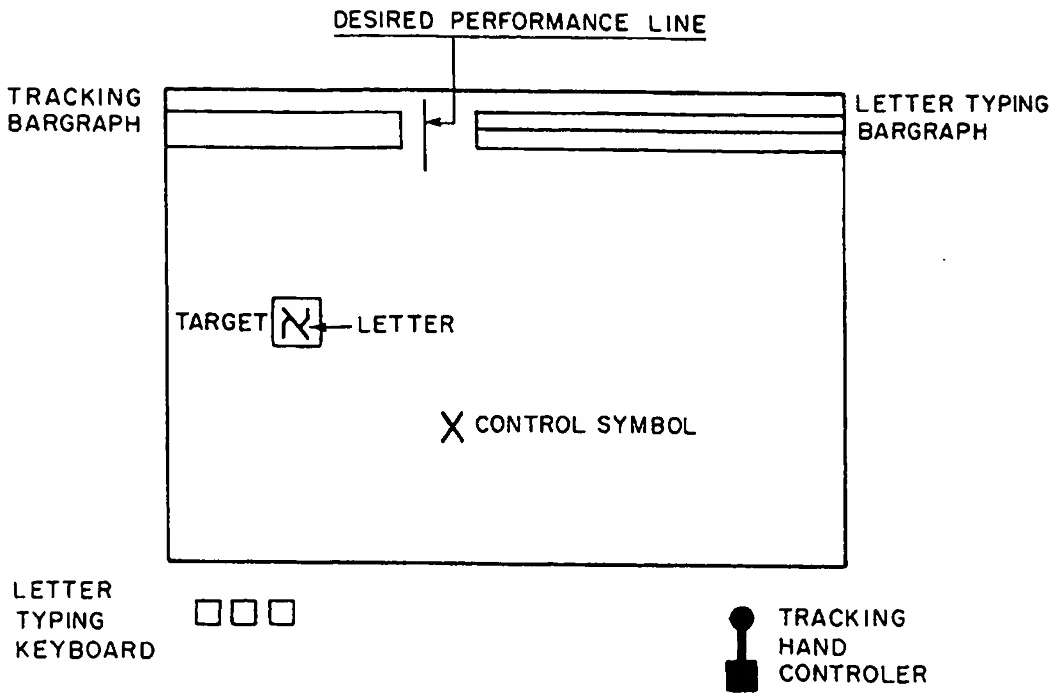

Figure 13.1 Subjects' display in concurrent performance of tracking and letter typing (from Brickner and Gopher 1981).

situations to optimize selected response strategies and enhance concurrent task performance.

A brief description of the techniques used to study these suggestions is in order before the main results of the experiments are presented. Figure 13.1 (Brickner and Gopher 1981) depicts one of our typical task situations (also see Gopher, Brickner, and Navon 1982; Gopher and Navon 1980; Navon et al. 1984). Subjects were presented with concurrent tracking and typing tasks. The tracking task required them to follow the movements of a computer-manipulated square by controlling the movement of an X symbol with a right-hand stick. For the typing task, subjects had to enter the corresponding chordic codes of letters presented within the tracking square, using a left-hand three-key keyboard. The upper portion of the screen displayed on-line dynamic feedback about their performance. Feedback indicators comprised a short vertical line representing desired performance levels, and two moving horizontal bar graphs representing actual performance levels. Desired performance was determined relative to normalized baseline performance distributions obtained for each subject. The distance between the bar graph of each task and the desired performance line reflected the momentary difference between actual and desired performance. This difference was computed continuously by subtracting the momentary error score (for tracking) and RT (for typing) from the desired score and dividing the outcome by the standard deviation of the subject's baseline distribution.

Task priorities were manipulated by moving the desired performance line from the center (equal priorities) to the left side (high priority for letter typing, low for tracking) or to the right side (high priority for tracking, low for typing). A priority level of, say, 70 percent for tracking corresponded to a level of performance that assumed the seventieth percentile in the baseline distribution of tracking performance by that subject. Instructions to put priority at 0.70 was actually a requirement to perform at a level better than the lowest 70 percent of the baseline performance levels. Other priority levels were determined in the same manner. Priority manipulations were complementary, so that when one task was increased, the other decreased, and vice versa.

Training under Variable Priorities

Brickner and Gopher (1981) investigated the effects of training under variable priorities in two dual-task experiments. In the first experiment, subjects performed concurrently the tracking and typing tasks presented in figure 13.1. In the second experiment, typing was paired with a discrete, self-paced, digit-classification task. Variable priority training was administered during two training sessions, comprising fifty three-minute trials, thirty-five of which fell under the relevant training conditions. Every trial was performed under one of five different relative priority levels (0.25, 0.35, 0.5, 0.65, 0.75).

Dual-task performance of the variable-priority group (VP) was compared with the performance of two other groups: an equal-priority (EP) group and a no-priority (NP) group. The EP group received the same augmented-feedback indicators as the VP group but performed all trials under the 0.5 equal-priority condition. The NP group received only verbal instructions with no augmented-feedback indicators. These subjects were instructed to consider both tasks as equally important, and they were given target performance levels similar to those of the EP group. Before the two training sessions, all groups had one session of practice under single-task conditions and dual-task conditions without feedback. Assignment of subjects to groups was determined at the end of this session.

The experimental results showed a large performance advantage for the variable-priority group over the other two groups. Moreover, the main differences were between the VP group and the other two groups that did not differ significantly from each other. Figures 13.2 and 13.3 depict the typing and tracking performance of the three groups on seven trials of the two training sessions. These were the trials in which the VP group performed under equal-priority instructions. For the EP and NP groups that performed only under equal-priority conditions, the plotted results are based on the trials that match in their sequential position those of the VP group. The advantage of the VP group is clear, despite the fact that unlike the other two groups, the equal-priority condition was practiced by this group in only 20 percent of the trials. The superiority of the VP group cannot be accounted for by the presence of augmented feedback. Performance levels of the EP group, which

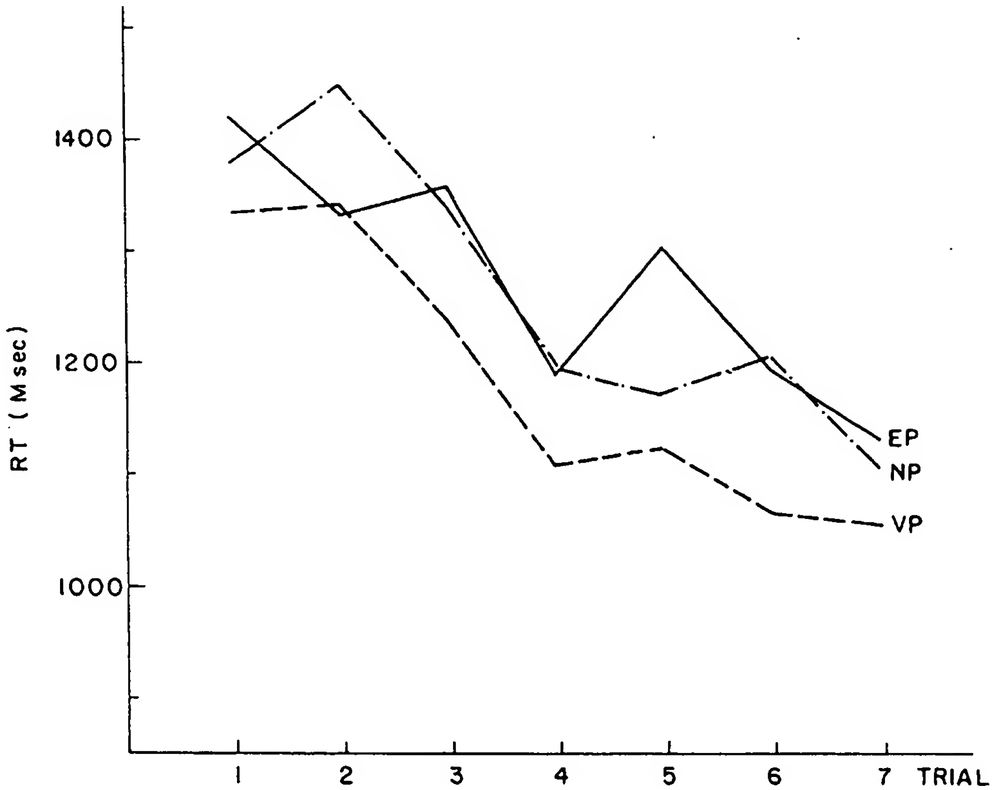

Figure 13.2 Letter-typing time (msec) during seven equal-priority trials, for each of three experimental groups (from Brickner and Gopher 1981).

was equally able to benefit from this feedback, were not significantly different from those of the NP group.

To further test the influence of variable-priority training, we gave subjects an additional (fourth) experimental session. In this session, the feedback indicators were removed from all groups. Instead, on each trial, the difficulty of the tracking and typing tasks was varied. There were three levels of difficulty for each task, and these levels were changed in a complementary manner; when the difficulty of typing increased, tracking difficulty decreased, and vice versa. There were four replications of the three commensurate difficulty combinations. Subjects were not given advance information about the change in difficulty. They were instructed to assign equal priority to tasks and to maintain a constant level of performance on both tasks during all trials. We reasoned that in order to maintain constant performance levels, subjects had to shift resources from performing the task that was made easy to performing the task whose difficulty was increased. Such adjustment had to be accomplished without prior information about the forthcoming changes, and augmented feedback was not available to help them supervise their performances.

The experimental data showed a clear superiority of the variable-priority group over the equal- and no-priority groups in maintaining a constant level of performance on both tracking and typing. Performance levels of the VP

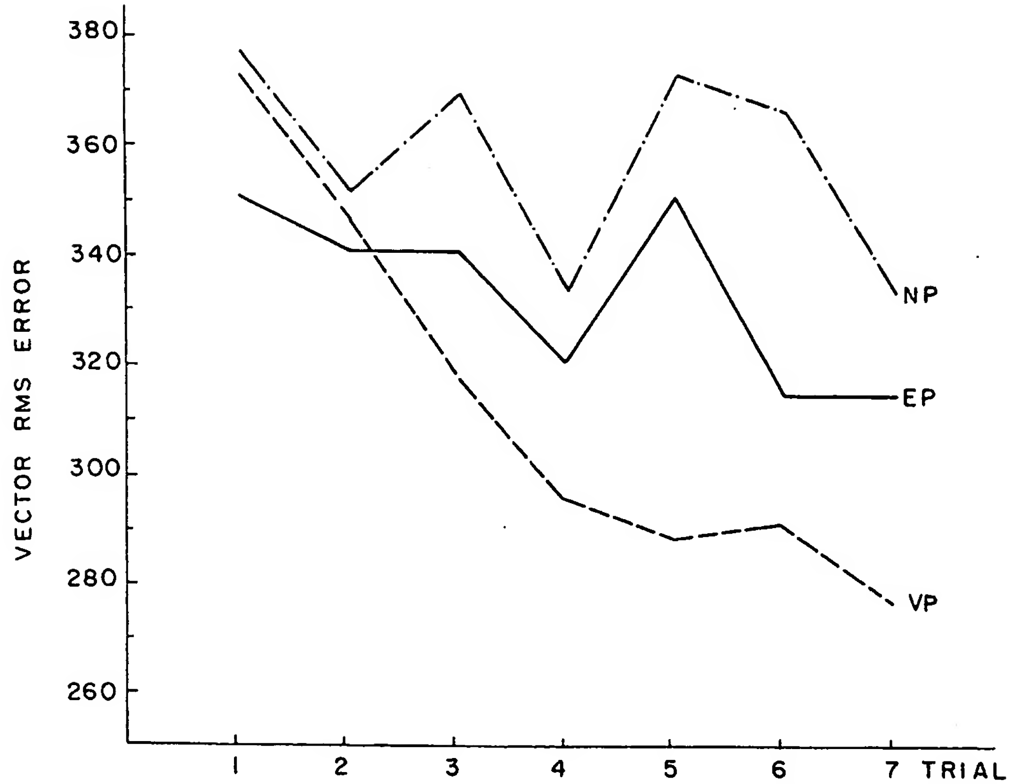

Figure 13.3 Tracking performance (RMS error) during seven equal-priority trials, for each of three experimental groups (from Brickner and Gopher 1981).

group were also much better than those of the other two groups. The EP and NP groups did not differ in their general performance levels, or in their ability to maintain performance despite the difficulty changes. The advantage of the VP group was smaller during the first two cycles of the new condition, but it gradually increased and stabilized over time (the Group × Replication interaction was highly significant).

Figure 13.4 depicts these results for the typing task; tracking errors had a very similar pattern. All groups were better and about equal in their single-task performance during this session, exhibiting similar performance slopes over the difficulty manipulation. It thus appears that the higher ability of the VP group in shifting efforts and maintaining performance under dual-task conditions does not stem simply from improved ability to perform each of the concurrent tasks. Rather, it stems from the control of behavior in the dual-task condition and the ability to detect changes and adjust efforts to cope with task demands.

A final test of the influence of variable-priority training was conducted in a fifth session. Subjects were required to pair one of the already-trained tasks with a new task of digit classification. Half of each of the original groups performed this task with tracking, the other half with typing. Additional superiority of VP training was revealed for the combination of tracking and

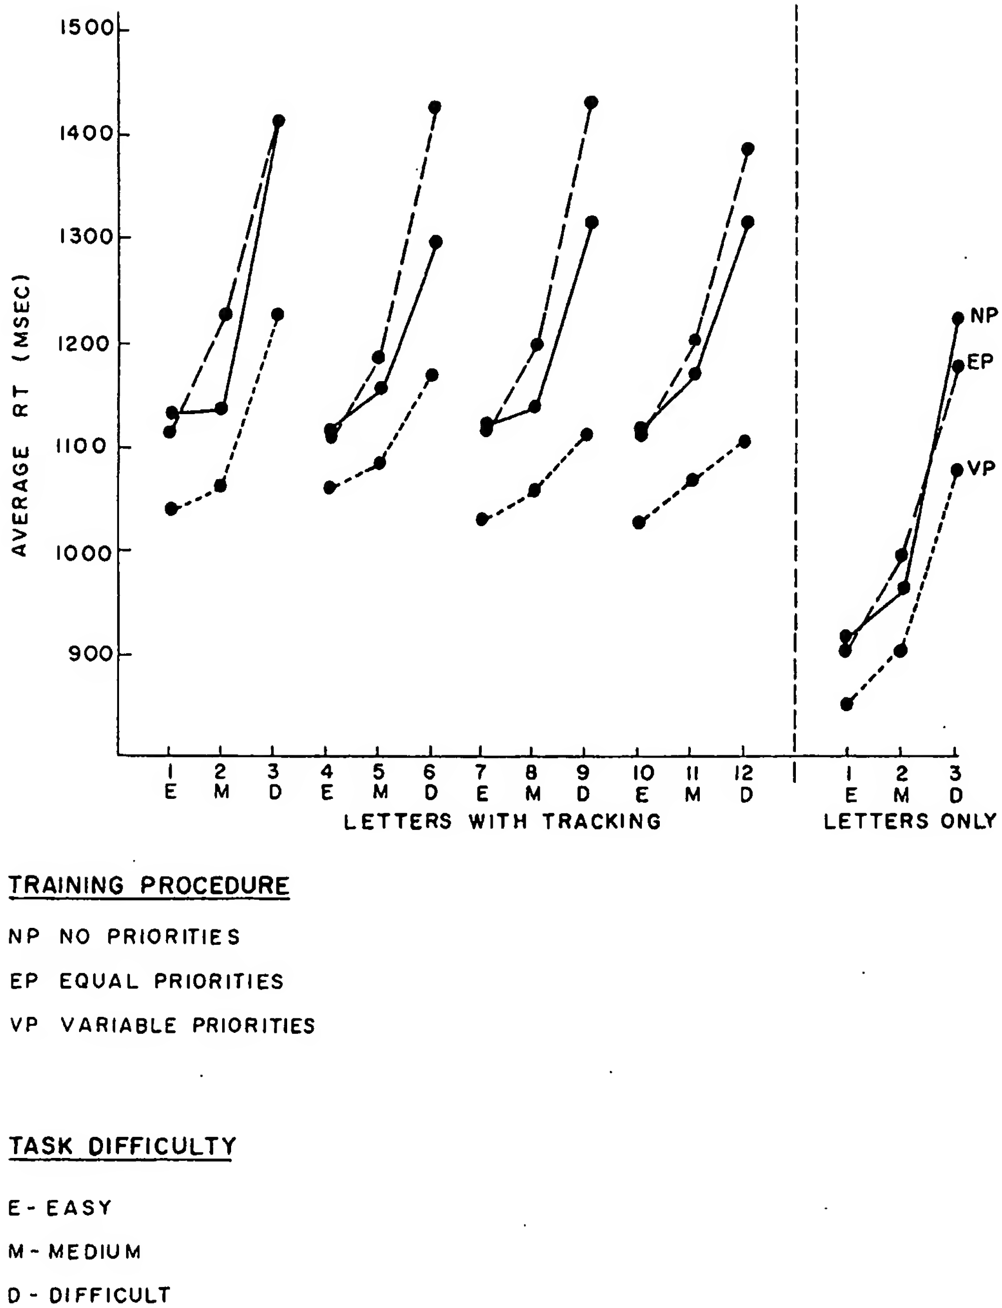

Figure 13.4 Four replications of dual-task letter-typing time (msec) and average single-task performance levels of the three training groups, under varying difficulty conditions (from Brickner and Gopher 1981).

digit classification, but not for the pairing of typing and digit classification (though there was a trend in the expected direction). One source of the difference between tracking and typing in this condition, as well as in the earlier training sessions, may be the better inherent feedback in the typing task. When typing, subjects receive immediate feedback about the accuracy of their responses. The feedback is much more indirect and obscure in the tracking task. Of the two classes of attention-control difficulties, tracking may have been deficient in both knowledge and control, while typing was deficient mainly in control.

The general pattern of results of the first experiment was replicated in a second experiment that investigated dual-task performance of typing and digit classification. A follow-up study by Spitz (1986, 1988) supported the findings of Brickner and Gopher (1981) with a new set of tasks. It also demonstrated the importance of the performance range over which priority changes are made.

Another aspect studied by Spitz was the influence of single-task training with variable priorities on dual-task performance. He reasoned that if part of the attention-control problem is the inability of subjects to establish graded standards of performance so as to release resources, then this ability should improve even with training under single-task conditions. He had two experimental groups practice in four single-task training sessions, using the same typing task as Brickner and Gopher (1981) and a different two-dimensional tracking task (size matching; cf. Navon et al. 1984).

A variable-priority (VP) group practiced each of the individual tasks under five different emphasis levels with augmented feedback. Another fixed-priority (FP) group practiced each task only under one emphasis level with the on-line augmented feedback. Following single-task training, subjects performed in a dual-task session of tracking and typing. During the first part of this session, they performed twelve three-minute dual-task trials in which they were instructed verbally to maximize their performance, assigning equal importance to both tasks. In the second part of the session, they were asked to maintain constant performance on both tasks, while undergoing the same sequence of subtle, complementary, difficulty changes as in Brickner and Gopher (1981).

Average performance levels at the end of single-task training showed no difference between the two experimental groups. In contrast, dual-task performance levels of subjects trained under variable emphasis were significantly better on both experimental tasks, as compared with fixed-priority training. The variable-priority group was also significantly better in coping with the difficulty-change conditions and in obeying the constant-performance instructions. A final comparison of importance in this context showed that despite the magnitude and significance of these effects, the absolute dual-task performance levels and maintenance ability after single-task variable-priority training were significantly lower than those obtained by a group trained in the same tasks with variable priorities under dual-task conditions.

It can thus be concluded that under dual-task conditions, one part of the control problem stems from the inability to establish, supervise, and maintain

graded control levels on each of the component tasks. Another part stems from the concurrency requirements, which account for the larger impact of variable-priority training under dual-task as compared with single-task conditions.

The knowledge and abilities that were acquired through variable-priority training have been shown to be present, influence, and benefit performance beyond the specific situation in which they are acquired. They affected performance without feedback indicators, when subtle implicit allocation adjustments were required, and when a new concurrent task was introduced. Such capabilities are the building blocks of attention strategies. They are the elements of an independent skill that can generalize and transfer beyond the specific situation in which it was originally trained.

Training through Emphasis Change on Subcomponents of Complex Tasks

During the last five years, we have continued to study the development of attention strategies in a more complex and dynamic situation, much closer to the task environments of driving, flying, and playing basketball discussed in section 13.1. The experiments at the Technion started as part of an international research collaboration conducted by twelve laboratories in four countries. It investigated the development of learning strategies to improve the behavior of trainees who practiced a highly complex and demanding computer game, which was designed jointly by the participating laboratories (Donchin, Fabiani, and Sanders 1989). The companion projects differed only in their proposed approach to training. Our project concentrated on the development of attention strategies to help performers cope with the high workload of the task.

We used an experimental game named Space Fortress, which was designed to simulate a complex and dynamic aviation environment. Players must control the movement of a spaceship on the display, fire missiles, and destroy a fortress. The fortress itself rotates, tracks, and fires at the spaceship. Other hostile elements are mines that appear every four seconds and actively chase the ship to damage it. There are two types of mines and different weapon systems to explode them. A large set of complex regulations specifies the rules of the game and the legal modes, intervals, and rates of responses. Another set of rules details point rewards and penalties. The subjects' main objective is to maximize their number of points during each three-minute game.

Figure 13.5 illustrates the game display. The game was carefully designed to include a representative sample of demand components. It combined a difficult, continuous manual-control task, several discrete and timed motor responses, visual monitoring, short-term and long-term memory load, dynamic division of attention, decision making, and resource management (Logie et al. 1989; Mane and Donchin 1989). Though comprising many dynamic elements and appearing cartoonlike to the naive eye, the game provided a rich and

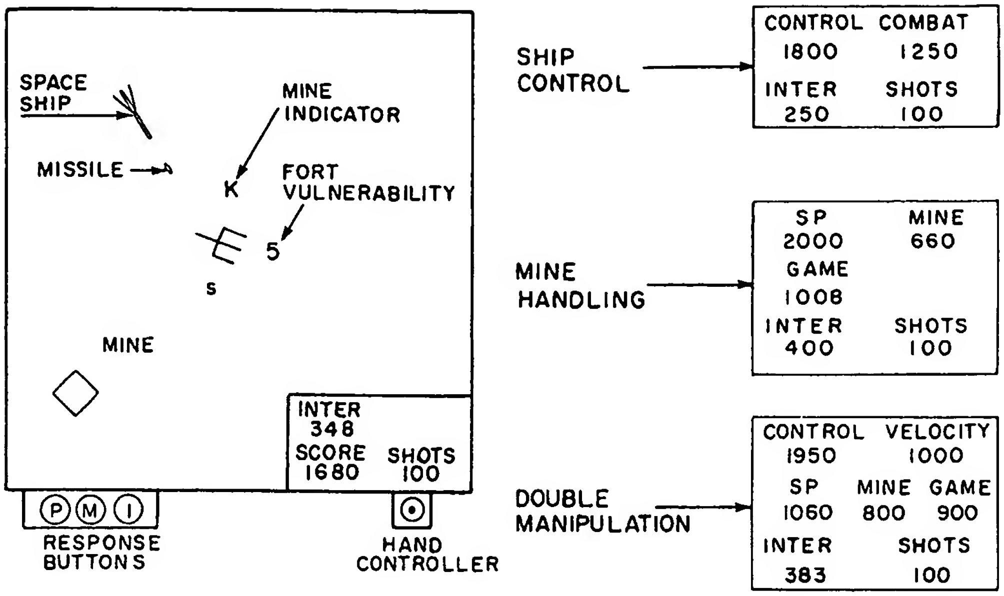

Figure 13.5 Subjects' display for the Space Fortress game. The three boxes on the right depict the special counters that were added in each of the emphasis manipulations. They replaced the information box presented for the standard game (from Gopher, Weil, and Siegal 1989).

highly controlled experimental environment in which parameters and manipulations were well specified.

Panic was the dominant first reaction of novice performers introduced to the game. Typical complaints were that "demands are too high; things happen too fast; too many events occur simultaneously; the situation is out of control." Performance scores during the first training session were all high negatives. However, with continuing experience, each control-group subject (who played the game with no special training) converged on a strategy that allowed him or her to survive and progress. The strategies differed among subjects. Still, given the algorithms and heuristics of the game, and the basic ability of subjects, the vast majority of them selected suboptimal strategies. Moreover, once a strategy was selected, it was never changed. After extensive practice, subjects just became more persistent and consistent in the application of their selected strategy. They progressed, but only within the limits and constraints that this strategy imposed (Foss et al. 1989; Gopher, Weil, and Siegal 1989).

This is yet another and more realistic manifestation of an attention-control and management problem. It emerged in the early phases of practice and continued to influence subjects' performance all the way through advanced stages of proficiency. Given the difficulty of the task, the emotional stress, and the lack of experience and skill, all of which typify the early stages of performance, it is hardly surprising that selected strategies were suboptimal. Less expected is the fact that new strategies were never explored.

Our approach to training in this difficult situation was based on leading subjects through a sequence of emphasis changes for subcomponents of the game. It embodied a direct outgrowth of the variable-priority method. Two

emphasis-change manipulations were developed, following a pilot study and task analysis. One focused on the control of the ship, the other on the handling of mines (Gopher, Weil, and Siegal 1989).

When an emphasis manipulation was introduced, subjects were instructed to concentrate on the selected aspect of the task and respond to others only insofar as the selected element was not neglected. Specific feedback indicators and score counters were added to inform and reward subjects about the success of their emphasis change, but none of the physical parameters of the standard game were eliminated or changed. The only changes were in the instruction to focus attention and in the score value for different types of responses.

The emphasis changes can be construed to have created a different figure-ground relationship among game components, but they did not influence their physical properties. This approach represents a radical change from a traditional part-task training philosophy that has guided other participants in the project (e.g., Fredriksen and White 1989; Newell et al. 1989). Rather than decomposing the task into simpler elements and training subjects on isolated parts before transferring to the whole, our method created different loci of emphasis within the whole without decomposing it into parts.

Several premises provided the foundations for our approach. Through emphasis changes, subjects are guided to play the game from different perspectives. Consequently, they are forced to explore alternatives rather than to commit themselves early to a single strategy. They are led, for example, to explore the value and the costs of different strategies. Such exploration may improve their attention control, equip them with alternative response strategies to meet dynamic changes of demands, and help them achieve a better match between their basic ability and the parameters of the task. Equally important is the fact that while focusing on an emphasized task component, subjects can also evaluate their peripheral attention abilities and improve their minimal control levels. This important opportunity does not arise if parts are trained in isolation, as in the traditional part-training approach.

Four groups of subjects were trained on the Space Fortress game for ten one-hour sessions. One group served as a control and did not receive any special training. Two groups experienced a single emphasis manipulation. One was given the ship's control, the other the mine-handling manipulation. A fourth group received both manipulations in alternation. All groups experienced the standard game in the first session. Emphasis training was conducted during sessions 2 through 6. From sessions 7 to 10, all special counters were removed and all groups returned to play the standard game, in which all task components had equal importance.

Figure 13.6 presents the average game scores of the four groups during the ten sessions of training. While all groups progressed with practice, the double-manipulation group had a substantial advantage over all other groups. Single-emphasis training was better than no emphasis change (control condition). However, the effect of double manipulation was much more pronounced and larger than the additive contributions of each individual emphasis manipula-

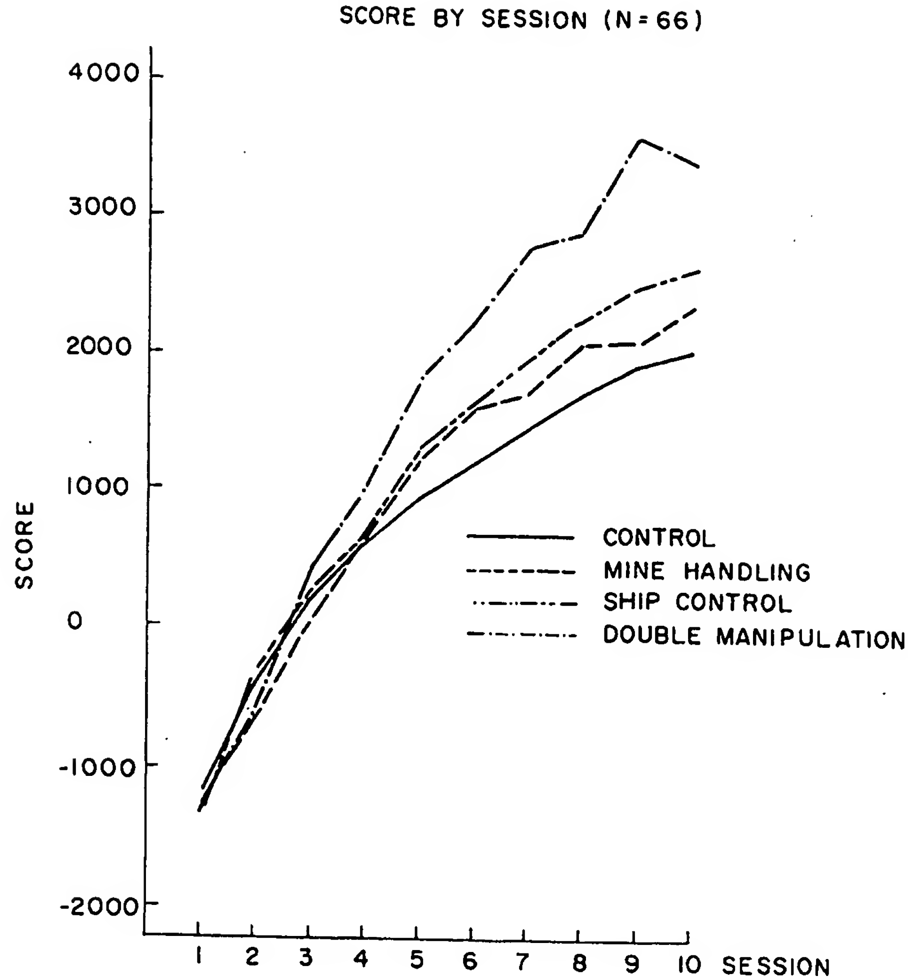

Figure 13.6 Learning curves of the four experimental groups. Performance is described in terms of total game scores. Special emphasis training was discontinued after session 6 (from Gopher, Weil, and Siegal 1989).

tion. Note also that there were no performance decrements when the special training was discontinued after the sixth session, and all auxiliary information was removed. The differences between groups continued to grow. It appears that by the seventh session subjects had already internalized their specialized knowledge and gained sufficient control to continue to improve on their own. Significant differences in the rates of progress indicated that the point of maximal difference between groups was yet to be reached.

Analysis of secondary performance indices demonstrated the differential influence of these emphasis manipulations. Thus, subjects were able to follow the instructions and shift their focus of attention. Although the overall scores of the two single-manipulation groups were not significantly different from each other, they were obtained by following different patterns of behavior. By comparison, the double-manipulation group was better able to progress in parallel on a wider range of game elements.

A second method of training that distinguished itself in the collaborative project was a hierarchical part-training approach developed by Fredriksen and White (1989). Based on an initial analysis of expert players in terms of game objectives, performance goals, and subgoals, they identified an inventory of twenty-seven part games. When practiced in isolation, part tasks were structured to lead subjects to gradually practice larger segments of the standard game. They were also given verbal tips and response recommedations derived from the analysis of expert behavior.

In a subsequent study, Fabiani et al. (1989) conducted an experiment to compare the relative effects of the emphasis change with the hierarchical part-training method. The experiment replicated the original ten-hour training schedules of the two methods but added five one-hour transfer sessions. During transfer, subjects were required to play the standard game while being given in each three-minute game trial one of seven secondary tasks. Secondary tasks were chosen from a battery developed in the original study by Logie et al. (1989) to investigate the load structure of the standard game. Thus, transfer sessions assessed the ability of subjects to cope with the additional load of a new concurrent task.

At the end of training, both groups were better than a control group that received no special instruction. Also, part-training subjects were significantly better than the emphasis-change group. In contrast, during transfer sessions, the effects were reversed. As during the variable priority experiments, subjects who trained under emphasis changes were better at coping with the additional secondary load, maintaining their original level of performance on the standard game and responding to secondary-task demands. Performance by the part-training group deteriorated markedly, declining on some task pairs to the control-group performance levels. The transfer superiority of emphasis-change training appeared both in proportional and absolute terms. It was largest on tasks imposing high working-memory load. A follow-up experiment showed that the advantage of the emphasis-training method persisted after fifteen additional hours of performing transfer tasks (Fabiani, Buckley, and Donchin, in press).

Recent research at the Technion employed a different set of secondary tasks and transfer conditions (Bareket 1990). This study generally replicated the results of Fabiani et al. (1989). Moreover, it showed that the advantage of the part-training method during the initial training period, which was observed in the Fabiani et al. study, completely disappeared when subjects in the emphasis-change group were given the same extra verbal tips.

One of the transfer conditions used by Bareket (1990) manipulated only the control dynamics of the ship without adding a secondary task. The control was changed from velocity to acceleration control, which required a radical change in subjects' response style. Subjects trained under emphasis change were significantly better in responding to this requirement than was the part-training group. It is clear from both studies (Fabiani et al. 1989; Bareket 1990)

that emphasis-change training under high-load conditions not only improved the ability of trainees to handle the specific task requirements, but also prepared them more generally to handle additional and new demands, and to reorganize an already acquired skill. These cases all imply reconfiguration of attention-control strategies.

Finally, perhaps the most daring test of our hypotheses has been conducted over the last three years at the flight-training school of the Israeli Airforce (Gopher et al. 1988). Problems of attention division, attention control, and coping with the load of a flight task have been a frequently cited cause of failure and washout from flight training. To address this difficulty, we gave a group of thirty-three flight cadets ten hours of emphasis-change training on a slightly modified version of Space Fortress. These modifications followed an analysis of the flight-training program to improve the face validity and organization of the game display. Game training was administered after the completion of a four-month period of ground school and flight training on a light aircraft, but before moving to a high-performance jet trainer.

The effects of game training were evaluated in eight subsequent test flights. They were compared with the flight performance of a matched control group of flight cadets having no game experience. The eight flights were part of the regular training program and a major decision point. At the end of this period, trainees either dropped out or were assigned to specialty programs (fighter, helicopter, or transport planes). Control and experimental subjects were matched by school psychologists based on their initial-selection battery scores and on flight performance in the light aircraft. The experimenters did not know subjects' initial ability scores. The whole project was presented to the subjects and instructor pilots as a study of individual differences, directed at developing a computerized test battery for flight selection. The experimental group comprised only a small part of two consecutive classes, but the flight instructors completed special evaluation forms for all trainees in both classes.

Flight performance scores showed large and significant advantages for trainees who received emphasis-change game training. As in the laboratory experiments, the differences between experimental and control groups increased across flights and were largest on the eighth flight. The average flight score of the experimental group was increased from 6 to 7, which represents a 30 percent increase in the probability of program completion. The advantage of game-trained subjects was largest in high-load maneuvers that required integration of several elements. A year and a half later, the actual percentage of graduates in the experimental group was twice as high as in the control group.

When considering the significance of these results, it is important to recognize certain key points. Although the game and flight situations have been construed to impose similar demands and to promote the development of similar skill components, the physical fidelity of the game is very low, and the general load of flight is higher by an order of magnitude. The game is not a flight simulator. At best, it constitutes a specific skill trainer. Nonetheless, the lessons learned and the skills acquired while playing the game helped trainees

to cope with the demands of flight, influencing their long-term approach to the acquisition of this complex and demanding skill.

13.4 CONCLUDING REMARKS

Our experiments have shown that while the potential to control attention and develop attention strategies exists, there are many difficulties and failures in realizing this potential when subjects have to cope with high-load, attention-demanding tasks. Nevertheless, through directed experience with different attention strategies, both task performance and attention control improve. Furthermore, these acquired capabilities and execution skills can be generalized; they may transfer to new situations and task demands, influencing the behavior of performers over extended time periods.

While ample evidence documents the effects and trainability of attention-management skills, the basic theoretical issues and empirical questions about them remained unresolved. We have proposed a working definition of attention strategies and a method for their development and manipulation, but this is a long way from uncovering the elements and structure of strategies, the modes in which they are represented in memory, and the linkage between them and the other skill components of a given task. In an earlier paper (Gopher, Weil, and Siegel 1989), we proposed that the response schemas associated with proficient behavior are the automated "hard-wired" products of early attention strategies.

To follow this line of thinking, attention strategies, previously defined as vectors of differential attention biases, may be alternately represented as vectors that combine performance objectives for elements of a complex task, in the service of a higher-level goal. Structurally, they can be construed as a set of production rules with attached gains and probability values binding together the elements of a complex task for the attainment of a specific objective. This suggestion may provide a framework with which to describe the process of moving from controlled strategies to automated schemas, consistent with contemporary views of skill training (Anderson 1981). However, even if the present view is accepted, we still lack knowledge about the building blocks and mechanisms that represent strategies and about the forces that drive them.

One approach toward closing this gap would be to further investigate what special propensity is acquired and generalized when subjects are trained under variable rather than fixed-priority conditions. Our results clearly demonstrate that the benifit of variable-priority training yields more than just an improved ability to perform the elements of a trained task. There is an additional and more general ability factor that develops independently of the specific factors governing task performance. Such capability may involve higher-level coordination and control functions that execute strategies and supervise behavior (e.g., Kahneman 1973; Logan 1985). Another possibility is that variable-priority training influences the form in which skill elements are organized and

represented in memory, allowing greater flexibility in the incorporation of new demands or reconfiguration of old elements.

Our present state of knowledge does not yet favor one of these alternatives over the other. However, they may be differentiated further through a better understanding of the context in which the acquired skill factors generalize. The success of the computer game training in flight school attests to the crucial importance of this research. We have already noted the low physical fidelity between the game trainer and the flight task, and the large gap between the two in complexity and task demands. What then could have been the commonalities that paved the way for their inclusion in our studies?

Part of the answer to this question may stem from the fact that both airplane piloting and our computer game training involved a difficult manual-control component, several discrete motor responses, visual orientation and scanning, long- and short-term memory load, decision making, and resource management, all interacting dynamically under high time pressure. These components are based on an analysis in terms of crude and generic constructs adopted from the information-processing literature (e.g., Gopher and Sanders 1984; Wickens 1984). Nonetheless, this framework may provide sufficient clarity and constraints to enable a valid comparative analysis of task demands. If so, then the prospects for human-performance models based on these constructs are brighter than expected. This framework may also define a viable manipulation environment for studying the architecture and generality of attention-control skills.

With an eye on growing needs for theoretical efforts to model the processing and response limitations of the human processing system, our experiments underscore the importance of considering attention-control problems as a new and major contributing source of performance decrements under time-sharing conditions. Previous models of attention control have contrasted data and resource limitations (Norman and Bobrow 1975), and resource scarcity versus output conflict (Navon and Gopher 1979). Some performance deficits previously associated with one or more of these factors should be reexamined to assess the likelihood of deficient control. The strong influence of training on the control of attention should also concern us from a methodological viewpoint. Instructions to subjects and experimental procedures should take into consideration more carefully the possible influence of attention strategies and control problems.

Finally, the relationship between attention control and consciousness requires examination. We construe attention strategies to be driven by a voluntary intent, but how many of these strategies, bits of knowledge, and control efforts are open to awareness or supervised consciously? The study of attention strategies along the lines proposed in this chapter may constitute a new and powerful vehicle for the study of the long-sought, but yet missing, formal link between consciousness and behavior. In contemporary scientific zeitgeist, so heavily biased towards connectionist and neural network modeling of bottom-up, emerging, and self-organizing patterns of behavior (Rumelhart and McClelland 1986; McClelland and Rumelhart 1986), a systematic study of the

power, properties, and operation rules of the top-down central executive may balance our prespective. Despite the fact that the notion of a central executive has a key position in many contemporary models of the human processing system, there has been little direct research on it. The present line of studies is a modest contribution in this direction.

NOTE

The preparation of this manuscript and our recent work in this area were supported by a grant from NASA, Ames Research Center.

REFERENCES

Allport, D. A., Antonis, B., and Reynolds, P. (1972). On the division of attention: A disproof of the single-channel hypothesis. *Journal of Experimental Psychology, 24,* 225–235.

Allport, D. A., and Styles, E. A. (1990). Multiple executive functions, multiple resources? Experiments in shifting attentional control of tasks. Unpublished manuscript.

Anderson, J. R. (1981). *Cognitive skills and their acquisition.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Bareket, T. (1990). The comparative effects of different training methods on the acquisition and transfer of complex skills. Master's thesis, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.

Brickner, M., and Gopher, D. (1981). Improving time-sharing performance by enhancing voluntary control on processing resources. Research Report, Technion, Research Centre for Work Safety and Human Engineering, HEIS-81-3, February.

Castiliano, U., and Umiltà, C. (1990). Size of attentional focus and efficiency of processing. *Acta Psychologica, 73,* 195–209.

Donchin, E., Fabiani, M., and Sanders, A. (Eds.) (1989). The learning strategies program: An examination of the strategies of skill acquisition. *Acta Psychologica, 71,* 1–312.

Egeth, H. (1977). Attention and preattention. In G. H. Bower (Ed.). *The Psychology of learning and motivation,* Vol. 11, 277–320.

Ericksen, C. W., and St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception and Psychophysics, 40,* 225–240.

Fabiani, M., Buckley, J., Gratton, G., Coles, M., and Donchin, E. (1989). The training of complex task performance. *Acta Psychologica, 71,* 259–300.

Fabiani M., Buckley, J., and Donchin, E. (N.d.). The effect of extended training on the performance of complex tasks. In press.

Foss, M. A., Fabiani, M., Mane, A., and Donchin, E. (1989). Unsupervised practice: The performance of the control group. *Acta Psychologica, 71,* 23–52.

Fredriksen, J. R., and White, B. Y. (1989). An approach to training based upon principal-task decomposition. *Acta Psychologica, 71,* 89–146.

Gopher, D. (1981). Performance tradeoffs under time-sharing conditions: The ability of human operators to release resources by lowering their standards of performance. In *Proceedings of the 1981 International Conference on Cybernetics and Society,* 609–613. Atlanta, GA.

Gopher, D. (1982). A selective attention test as a predictor of success in flight training. *Human Factors, 24,* 173–183.

Gopher, D. (1990). Dual-tasks. In N. Eysenk (Ed.), *The Blackwell Dictionary of Cognitive Psychology*, 23–28. Oxford: Basil-Blackwell.

Gopher, D., and Braune, R. (1984). On the psychophysics of workload: Why bother with subjective measures. *Human Factors, 26*, 519–532.

Gopher, D., Brickner, M., and Navon, D. (1982). Different difficulty manipulations interact differently with task emphasis: Evidence for multiple resources. *Journal of Experimental Psychology: Human Perception and Performance, 8*, 146–157.

Gopher, D., and Donchin, E. (1986). Workload: An examination of the concept. In K. Boff and L. Kaufman (Eds.), *Handbook of human perception and performance*, 1–49. New York: John Wiley.

Gopher, D., and Kahneman, D. (1971). Individual differences in attention and the prediction of flight criteria. *Perceptual and Motor Skills, 33*, 1335–1342.

Gopher, D., and Navon, D. (1980). How is performance limited: Testing the notion of central capacity. *Acta Psychologica, 46*, 161–180.

Gopher, D., and North, R. N. (1977). Manipulating the conditions of training in time-sharing performance. *Human Factors, 19*, 583–593.

Gopher, D., and Sanders, A. F. (1984). S-Oh-R: Oh stages! Oh resources. In W. Printz and A. F. Sanders (Eds.), *Cognition and motor processes*, 231–254. Berlin: Springer-Verlag.

Gopher, D., Weil, M., Bareket, T., and Caspi, S. (1988). Fidelity of task structure as a guiding principle in the development of skill trainers based upon complex computer games. In *Proceedings of the 32nd Annual Meeting of the Human Factors Society*, 1266–1270. Anaheim, CA.

Gopher, D., Weil, M., and Siegel, D. (1989). Practice under changing priorities: An approach to training of complex skills. *Acta Psychologica, 71*, 147–179.

Hirst, W., Spelke, E. S., Reaves, C. C., Caharack, G., and Neisser, U. (1980). Dividing attention without alternation or automaticity. *Journal of Experimental Psychology: General, 109*, 98–117.

Jonides, J. (1983). Further towards a model of the mind's eye-movements. *Bulletin of The Psychonomic Society, 21*, 247–250.

Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice-Hall.

Kahneman, D., Ben Ishai, R., and Lotan, M. (1973). Relation of a test of attention to road accidents. *Journal of Applied Psychology, 58*, 113–115.

Kahneman, D., and Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman and D. Davis (Eds.), *Varieties of attention*, 29–63. Orlando, FL: Academic Press.

Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica, 60*, 193–210.

Logie, R., Baddeley, A., Mane, A., Donchin, E., and Sheptak, R. (1989). Working memory in the acquisition of complex cognitive skills. *Acta Psychologica, 71*, 53–88.

Mane, A., and Donchin, E. (1989). The space fortress game. *Acta Psychologica, 71*, 17–22.

McClelland, J. L., and Rumelhart, D. E., (Eds). (1986). *Parallel Distributed Processing: Explorations in the microstructure of cognition. Vol. 2: Applications*. Cambridge, MA: Bradford Books/MIT Press.

Navon, D., and Gopher, D. (1979). On the economy of the human processing system. *Psychological Review, 86*, 214–253.

Navon, D., Gopher, D., Chillag, M., and Spitz, G. (1984). On separability of and interference between tracking dimensions in dual-axis tracking. *Journal of Motor Behavior, 16*, 364–392.

Newell, K. M., Carlton, M. J., Fisher, A. T., and Rutter, B. G. (1989). Whole-part training in the acquisition of a complex perceptual-motor skill. *Acta Psychologica, 71*, 197–216.

Norman, D. A., and Bobrow, D. J. (1975). On data limited and resource limited processes. *Cognitive Psychology, 7*, 44–64.

Norman, D. A., and Shallice, T. (1981). Attention to action: Willed and automatic control of behavior. In M. Lansman and E. Hunt (Eds.), *Proceedings of the Lake Wilderness Attention Conference*. N-25, Department of Psychology, University of Washington, Seattle, WA.

North, R. A., and Gopher, D. (1976). Measures of attention as predictors of flight performance, *Human Factors, 18*, 1–14.

Posner, M. I., and Rothbart, M. K. (1986). The concept of energy in psychological theory. In R. J. Hockey, A. W. Gaillard, and M. G. H. Coles (Eds.). *Energetics and human information processing*, 23–42. *Dordrecht: Martinus* Nijhoff Publishers.

Posner, M. I., Snyder, C. R. R., and Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General, 109*, 160–174.

Rumelhart, D. E., and McClelland, J. L. (Eds.) (1986). *Parallel Distributed Processing: Explorations in the microstructure of cognition. Vol. I: Foundations*. Cambridge. MA: Bradford Books/MIT Press.

Schneider, W., and Fisk, A. D. (1982). Dual task automatic and control processing: Can it be done without cost? *Journal of Experimental psychology: Learning, Memory, and Cognition, 8*, 261–278.

Schneider, W., and Fisk, A. D. (1984). Automatic category search and its transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 1–13.

Sperling, G., and Melchner, M. J. (1978). The attention operating characteristic: Some examples from visual search. *Science, 202*, 315–318.

Spitz, G. (1986). Voluntary control of mental resource allocation. Ph.D. diss., Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.

Spitz, G. (1988). Flexibility in resource allocation and the performance of time-sharing tasks. (1988). *Proceedings of the 32nd Annual Meeting of the Human Factors Society*, 1466–1469. Anaheim, CA.

Strayer, D. L., and Kramer, A. F. (1990). Attentional requirements of automatic and controlled processing. *Journal of Experimental Psychology: Learning, Memory and Cognition, 16*, 291–304.

Wickens, C. D. (1980). The structure of attentional resources. In R. S. Nickerson (Ed.), *Attention and Performance VIII*, 239–257. Hillsdale, NJ: Lawrence Erlbaum Associates.

Wickens, C. D. (1984). Processing resources in attention. In R. Parasuraman and D. R. Davis (Eds.), *Varieties of Attention*. Orlando, FL: Academic Press.

Wickens, C. D., and Gopher, D. (1977). Control theory measures of tracking as indices attention allocation strategies. *Human Factors, 19*, 349–365.

Ye, Y. Y., and Wickens, C. D. (1988). Dissociation of performance and subjective measures of workload. *Human Factors, 30*, 111–120.

What causes transfer to be positive, negative, or zero? To oversimplify the research results somewhat, it is apparent that positive transfer occurs when the training program and the target task are quite similar (in fact, if they are identical, transfer is usually about as positive as it can be, although there are some exceptions). Extreme differences between the training and target task typically produce zero transfer. Learning to type, for example, does not help learning to swim or drive an automobile. Negative transfer occurs from a particularly unique set of circumstances relating to perceptual and response aspects of the task, which we will describe later. We will first consider the goal of the training system design, which is to produce maximum positive transfer by increasing the similarity of the training device to the target task. The issue of *how* similarity should be increased defines training system fidelity. Next we will consider negative transfer between old and new tasks, and finally the role of analogies in transfer.

Training System Fidelity We stated that maximum positive transfer would generally occur if all elements of a task were identical to the target task. Does this mean that training simulators should resemble the real world as closely as possible? In fact, the answer to this question is no for a number of reasons (Hopkins, 1975; Jones, Hennessy, & Deutsch, 1985; Schneider, 1985). First, highly realistic simulators tend to be very expensive, but their added realism may add little to their TER (Hawkins, 1987). Second, in some cases, high similarity, if it does not achieve complete identity with the target environment, may be detrimental by leading to incompatible response tendencies or strategies. For example, there is little evidence that motion in flight simulators, which cannot approach the actual motion of the aircraft, offers any positive transfer benefits (Hawkins, 1987; Lintern, 1987). Finally, if high realism presents overwhelming complexity, it may so increase workload and divert attention from the critical skill to be learned that effective learning is inhibited.

Instead of total fidelity in training, researchers have emphasized understanding *which* components of training should be made similar to the target task and which may be less important (Holding, 1987; Singley & Andersen, 1989). For example, training simulators for a sequence of procedures may be of very low fidelity to the target task, as long as the sequences of steps are compatible (Hawkins, 1987; Holding, 1987). There is also evidence that a major portion of learning complex skills may be tied to the recognition and use of perceptual consistencies, or *invariants*, in the environment (Lintern, Roscoe, & Sivier, 1990; Schneider, 1985), a concept discussed in Chapter 4. Hence it is important that simulators be designed so that the useful consistencies in the target task are preserved and made visible, and the learner is made aware of them in the training environment (Lintern, Thomley-Yates, Nelson, & Roscoe, 1987). For example, beginning drivers need to recognize the important features of the highway (the heading of the vehicle relative to the vanishing point, more than the momentary deviation) that allow them to steer the vehicle effectively. A training simulator should incorporate this information clearly.

In summary, it appears that some types of departures from full fidelity do

not have the detrimental impact on transfer that would be predicted from a straightforward view that maximum similarity produces maximum transfer. Furthermore, some training strategies are based on the assumption that carefully planned departures from similarity can actually enhance transfer if they focus the attention of trainees on critical components of the task.

Negative Transfer The issue of negative transfer is a critical one, as the continued emergence of new technology and different system designs leads operators to switch from one system to another. What causes the skills acquired in one setting to inhibit some aspect of performance in a different one? A history of research in this area (Holding, 1976; Martin, 1965; Osgood, 1949) seems to reveal that the critical conditions for negative transfer are related to the stages of processing. When the two situations have highly similar (or identical) stimulus elements but different response or strategic components, transfer will be negative, particularly if the new and old responses are incompatible with one another (i.e., the new and the old response cannot easily be given at the same time). Stated another way, Singley and Andersen (1989) have argued that a major source of negative transfer is from the positive transfer of *inappropriate* responses. The relationship between the similarity of stimulus and response elements and transfer is shown in Table 6.1.

It is important to realize that many real-world tasks involve the transfer of a large number of different components, and generally more of them produce positive rather than negative transfer. Hence, most transfers from one similar task to another are generally positive. However, the critical design questions may be focused on those aspects of the difference between training and transfer (or between an old and new system) that *do* involve incompatible responses or inappropriate strategies. For example, consider two word-processing systems that present identical screen layouts but require a different set of key presses to accomplish the same editing commands. A high level of skill acquired through extensive training on the first system will show some interference with transfer to the second, even though the overall transfer will be positive. Or consider two control panels that both require a lever movement to accomplish a function. However, in one panel the lever must be pushed up, and in the other it must be pushed down.

Table 6.1 RELATIONSHIP BETWEEN OLD AND NEW TASK

Stimulus Elements	Response Elements	Transfer
Same	Same	++
Same	Difference	−
	(incompatible)	−−
Different	Same	+
Different	Different	0

Journal of Experimental Psychology:
Human Perception and Performance
2001, Vol. 27, No. 4, 763–797

0096-1523/01/$5.00 DOI: 10.1037//0096-1523.27.4.763

Executive Control of Cognitive Processes in Task Switching

Joshua S. Rubinstein
Federal Aviation Administration

David E. Meyer and Jeffrey E. Evans
University of Michigan

In 4 experiments, participants alternated between different tasks or performed the same task repeatedly. The tasks for 2 of the experiments required responding to geometric objects in terms of alternative classification rules, and the tasks for the other 2 experiments required solving arithmetic problems in terms of alternative numerical operations. Performance was measured as a function of whether the tasks were familiar or unfamiliar, the rules were simple or complex, and visual cues were present or absent about which tasks should be performed. Task alternation yielded switching-time costs that increased with rule complexity but decreased with task cuing. These factor effects were additive, supporting a model of executive control that has goal-shifting and rule-activation stages for task switching. It appears that rule activation takes more time for switching from familiar to unfamiliar tasks than for switching in the opposite direction.

Traditionally, experimental psychology has focused on studying repetitive performance of individual perceptual–motor and cognitive tasks. Nevertheless, daily life often requires performing multiple tasks either simultaneously or in rapid alternation, as when people prepare meals while tending children or drive automobiles while operating cellular telephones. To explain how such multiple-task performance is achieved, some theorists have proposed that executive control processes supervise the selection, initiation, execution, and termination of each task (e.g., Baddeley, 1986; Duncan, 1986; Logan, 1985; Meyer & Kieras, 1997a, 1997b; Norman & Shallice, 1986; Shiffrin & Schneider, 1977). These proposals extend classical ideas about voluntary willed action (James, 1890), which may be elaborated in terms of concepts from computer science and multitasking operating systems (Kieras, Meyer, Ballas, & Lauber, 2000; Neisser, 1967).

Given this state of affairs, research on executive mental control requires asking various detailed analytical questions (Monsell, 1996). Are executive control processes really separable from the basic perceptual–motor and cognitive processes used for performing individual tasks? How might executive control processes establish priorities among individual tasks and allocate resources to them during multiple-task performance? Of what functionally distinct subcomponents do executive control processes consist? The present article provides further answers to such questions through experiments with a successive-tasks procedure developed and used previously for studying executive control processes that enable task switching (e.g., Allport, Styles, & Hsieh, 1994; Botwinick, Brinley, & Robbin, 1958; Dark, 1990; Garcia-Ogueta, 1993; Jersild, 1927; Keele & Hawkins, 1982; Los, 1996; Meiran, 1996; Meiselman, 1974; Rogers & Monsell, 1995; Spector & Biederman, 1976; Weber, Burt, & Noll, 1986).

In subsequent sections of this article, we start by briefly reviewing available theories of executive control processes. Some past studies of task switching whose results bear on the veracity of these theories are summarized next. Then a new model of executive mental control in task switching is introduced. To test this model and to demonstrate its potential heuristic value, we report four experiments with the successive-tasks procedure. On the basis of data from them, we propose that task switching entails at least two functionally distinct stages of executive control, goal shifting and rule activation, which are separable from the basic perceptual–motor and cognitive processes used for performing individual tasks. Our proposed stage model provides coherent explanations of numerous previous findings about task switching and suggests promising directions for future research on executive mental control.

Joshua S. Rubinstein, William J. Hughes Technical Center, Federal Aviation Administration, Atlantic City, New Jersey; David E. Meyer, Department of Psychology, University of Michigan; Jeffrey E. Evans, Department of Rehabilitation Psychology and Neuropsychology, University of Michigan Medical Center.

Portions of this research are based on a doctoral dissertation submitted by Joshua S. Rubinstein to the Department of Psychology, University of Michigan (Rubinstein, 1993). Initial reports of the present results have been made at meetings of the Psychonomic Society (Rubinstein, Meyer, & Evans, 1994) and the Cognitive Neuroscience Society (Rubinstein, Evans, & Meyer, 1994). Funding was provided by National Institute of Mental Health Grant MH16892 to the University of Michigan.

We thank Henry Buchtel, Douglas Medin, Sarah Winans Newman, and Patricia Reuter-Lorenz for encouragement and suggestions about the work reported herein. Comments by Alan Allport, Tom Carr, David Kieras, Gordon Logan, Ulrich Mayr, Stephen Monsell, Nora Newcombe, Tim Shallice, and Saul Sternberg about previous versions of this article are also gratefully acknowledged.

Correspondence concerning this article should be addressed to Joshua S. Rubinstein, Building 315, William J. Hughes Technical Center, Federal Aviation Administration, Atlantic City International Airport, Atlantic City, New Jersey 08405, or to David E. Meyer, Cognition and Perception Program, Department of Psychology, University of Michigan, 525 East University, Ann Arbor, Michigan 48109-1109. Electronic mail may be sent to Josh.Rubinstein@tc.faa.gov or to demeyer@umich.edu.

Theories of Executive Control Processes

For now, we focus on three representative theories: the attention-to-action (ATA) model (Norman & Shallice, 1986), the frontal-lobe executive (FLE) model (Duncan, 1986), and the strategic response-deferment (SRD) model (Meyer & Kieras, 1997a, 1997b, 1999). These theories are especially relevant because they exemplify how task switching might be mediated by separable executive control processes that prepare systematically for transitions between successive tasks.

Attention-to-Action Model

The ATA model of Norman and Shallice (1986) has three subcomponents: *action schemas, contention scheduling*, and a *supervisory attentional system* (SAS).

Action schemas are specialized routines for performing individual tasks that involve well-learned perceptual–motor and cognitive skills. Each action schema has a current degree of activation that may be increased by either specific perceptual "trigger" stimuli or outputs from other related schemas. When its activation exceeds a preset threshold, an action schema may direct a person's behavior immediately and stereotypically toward performing some task. Moreover, on occasion, multiple schemas may be activated simultaneously by different trigger stimuli, creating error-prone conflicts if they entail mutually exclusive responses (e.g., typing on a keyboard and answering a telephone concurrently).

To help resolve such conflicts, the ATA model uses contention scheduling. It functions rapidly, automatically, and unconsciously through a network of lateral inhibitory connections among action schemas whose response outputs would interfere with each other (cf. Rumelhart & Norman, 1982). Through this network, an action schema (e.g., one for keyboard typing) that has relatively high current activation may suppress the activation of other potentially conflicting schemas (e.g., one for telephone answering). Contention scheduling allows task priorities and environmental cues to be assessed on a decentralized basis without explicit top-down executive control (Shallice, 1988). However, this may not always suffice to handle conflicts when new tasks, unusual task combinations, or complex behaviors are involved.

Consequently, the ATA model also has an SAS. The SAS guides behavior slowly, flexibly, and consciously in a top-down manner. It helps organize complex actions and perform novel tasks by selectively activating or inhibiting particular action schemas, superseding the cruder bottom-up influences of contention scheduling and better accommodating a person's overall capacities and goals. For example, one might expect the SAS to play a crucial role during switches between unfamiliar incompatible tasks that are not ordinarily performed together.

Depending on conditions that prevail during multiple-task performance, the ATA model accounts qualitatively for a variety of empirical phenomena. In particular, slips of action that occur during daily activities (e.g., Reason, 1990) may stem from temporary failures of the SAS to regulate contention scheduling adequately. SAS failures may also explain behavioral abnormalities in patients with frontal-lobe brain damage (Shallice, 1982, 1988, 1994).

Frontal-Lobe Executive Model

Assumptions similar to those of the ATA model have been embodied in the FLE model of Duncan (1986). It has three main components: *goal lists, means–ends analysis procedures*, and *action structures*. Goal lists represent a person's current set of prioritized intentions. Means–ends analysis, somewhat like the SAS (cf. Norman & Shallice, 1986), updates the contents and order of goals in working memory, taking account of how well they are being achieved over time. Supplementing such functions, the action structures of the FLE model constitute a large store of procedural knowledge for goal-directed behaviors embodied as sets of condition–action production rules (cf. Allport, 1980; J. R. Anderson, 1983, 1993; Hunt & Lansman, 1986; Logan, 1985; Newell, 1973, 1990; Townsend, 1986). The conditions of these rules refer to goals and perceptual stimuli; the actions involve responses to achieve the goals (e.g., IF THE GOAL IS TO DO TASK *A* AND THE STIMULUS IS *S*, THEN PRODUCE RESPONSE *R*). Action structures composed of such rules are functionally analogous to the ATA model's action schemas.

Furthermore, according to Duncan (1986), goal lists and means–ends analysis are implemented primarily in the brain's frontal lobes. The FLE model implies that damage to particular frontal-lobe regions may disrupt people's ability to maintain and pursue their goals, reducing their effectiveness in planning and performing multiple tasks. This implication also agrees with claims of some other theorists (e.g., Kimberg & Farah, 1993; Shallice, 1994; Stuss & Benson, 1986).

Strategic Response-Deferment Model

Additional detailed ideas about how executive control contributes to multiple-task performance have been provided by Meyer and Kieras (1997a, 1997b, 1999). Using a production-rule formalism, they constructed an executive-process interactive control (EPIC) architecture that combines various components of the human information-processing system in a unified theoretical framework (cf. J. R. Anderson, 1983, 1993; Card, Moran, & Newell, 1983; Newell, 1990). EPIC includes perceptual, cognitive, and motor processors interfaced with working-memory stores whereby multiple-task performance can be described computationally. For example, on the basis of EPIC, Meyer and Kieras proposed an SRD model that simulates performance in a traditional dual-task paradigm, the psychological refractory-period (PRP) procedure (Bertelson, 1966; Kantowitz, 1974; Pashler, 1994; Smith, 1967; Welford, 1952, 1959, 1967).

The PRP procedure exemplifies a *simultaneous-tasks procedure*. On each discrete trial of this procedure, a stimulus is presented for the first of two tasks that entail stages of processing such as stimulus identification, response selection, and movement production. In response to the Task 1 stimulus, a participant must react quickly and accurately. Soon after the Task 1 stimulus, another stimulus is presented for the second task, separated by a short (e.g., ≤ 1 s) stimulus-onset asynchrony (SOA). In response to the Task 2 stimulus, the participant must again react quickly and accurately. However, instructions for the PRP procedure require that Task 1 receive higher priority than Task 2 (e.g., Pashler, 1984; Pashler & Johnston, 1989), and reaction times (RTs) are measured

to assess the extent to which the two tasks interfere with each other.

To characterize this interference, the SRD model of Meyer and Kieras (1997a, 1997b, 1999) assumes that performance during the PRP procedure involves three sets of production rules. One rule set implements operations for Task 1 (e.g., selecting Task 1 responses). A second rule set implements operations for Task 2 (e.g., selecting Task 2 responses). A third executive-process rule set schedules these operations so that instructions about task priorities are obeyed and conflicts do not occur over the use of limited-capacity perceptual–motor components. By manipulating task goals and strategy notes in working memory, the executive process permits the Task 1 production rules to select and send Task 1 responses to an appropriate motor (e.g., manual or vocal) processor as soon as possible, regardless of the SOA. Task 2 production rules are also permitted to select Task 2 responses concurrently with Task 1 response selection (cf. Pashler, 1994; Welford, 1967). At short SOAs, however, the model's executive process defers movement production for Task 2 by storing selected Task 2 responses temporarily in working memory until Task 1 has been completed. This ensures that Task 2 responses do not inadvertently precede or interfere with Task 1 responses at peripheral levels. After completion of Task 1 on a trial, the executive process permits any previously selected and stored Task 2 response to be produced. Also, if Task 2 response selection has not yet started, then a subsequently selected Task 2 response is permitted to be produced immediately. Such temporal overlapping and interleaving of Task 1 and Task 2 processes accounts well for patterns of additive and interactive factor effects on empirical mean RTs from the PRP procedure (e.g., Hawkins, Rodriguez, & Reicher, 1979; Karlin & Kestenbaum, 1968; McCann & Johnston, 1992; Meyer et al., 1995; Pashler, 1990; Schumacher et al., 1999, 2001). In essence, the SRD model demonstrates how some basic ideas from the ATA and FLE models can be formalized and tested successfully against quantitative data.

Tentative Theoretical Hypotheses

Given the success of these models in accounting qualitatively and quantitatively for major phenomena associated with multiple-task performance, some interesting theoretical hypotheses may be advanced. Perhaps executive control processes really do exist, and perhaps they incorporate multiple separable subcomponents that enable task switching. Thus, subsequent sections of this article consider these hypotheses and ways of testing them further.

Task Switching and the Successive-Tasks Procedure

Some evidence about the existence and separability of component executive control processes comes from a *successive-tasks procedure* for studying task switching. The successive-tasks procedure is similar in certain respects to the PRP simultaneous-tasks procedure mentioned earlier. However, there are also conceptually important differences between these two procedures. In what follows, we discuss the successive-tasks procedure more fully, and we summarize representative results that have been obtained with it.

Successive-Tasks Procedure

Several basic features characterize the successive-tasks procedure (Monsell, 1996).

Assignment of task priorities. When the procedure is implemented, equal priorities are typically assigned to the individual tasks between which participants must switch. This assignment contrasts with that of the PRP simultaneous-tasks procedure, wherein one task is primary and the other secondary. Consequently, scheduling the stages of processing for the successive-tasks procedure may be relatively simple in certain respects, lessening the demands placed on executive mental control (cf. Kieras et al., 2000).

Temporal sequence of stimulus events. The assignment of equal task priorities is encouraged by the temporal sequence of stimulus events during the successive-tasks procedure. In this procedure, the stimulus for the next task is never presented until after a response to the current task stimulus has occurred. This constrains the *response–stimulus interval* (RSI) to be nonnegative and the SOAs to be all relatively long (i.e., equal to or greater than concomitant RTs). Thus, unlike the PRP simultaneous-tasks procedure, the successive-tasks procedure provides little, if any, opportunity to overlap the stages of processing for two or more tasks. Again, this may lessen the demands imposed on executive mental control.

Composition of stimulus–response mappings. Nevertheless, these demands can still be substantial because of the stimulus–response (S-R) mappings that are typically used during the successive-tasks procedure. Here the stimuli and responses are often the same for all of the tasks; one task's S-R mapping may differ from another's only in terms of which specific responses are associated with which specific stimuli. Consequently, under this procedure, task switching is potentially susceptible to proactive interference reminiscent of what occurs during verbal learning and memory (Allport et al., 1994; cf. Crowder, 1976). To cope with such interference, executive control processes may need to incorporate response monitoring and inhibitory mechanisms.[1]

Theoretical relevance. Because of its characteristic features, the successive-tasks procedure is especially relevant to addressing some important issues about the nature of multiple-task performance. It allows an investigator to examine how executive control processes enable task switching when task stimuli do not overlap temporally and responses to them need not be selected or produced in parallel, but alternative S-R mappings may induce considerable proactive interference between tasks. As discussed next, past studies conducted under such conditions have yielded many informative results (for other reviews of the literature, see Monsell, 1996; Monsell, Yeung, & Azuma, 2000).

[1] In contrast, the PRP simultaneous-tasks procedure often involves two tasks whose stimuli and responses involve different sensory and motor modalities, respectively. Given these differences, the S-R mappings for the primary and secondary tasks would be clearly distinct, there would be no perceptual or motor conflicts between them, and proactive interference would not prevail. This may encourage the use of sophisticated time-sharing algorithms and concurrent response-selection processes (Meyer & Kieras, 1997a, 1997b), whereas the successive-tasks procedure does not (Kieras et al., 2000).

Jersild's (1927) Study

An influential early study with a precursor of the successive-tasks procedure was conducted by Jersild (1927).[2] It yielded substantial time costs of task switching whose magnitudes depended on the complexity of the operations that were performed during each task. This dependence could bear on the nature of underlying executive control processes.

In one experiment, Jersild gave participants columns of two-digit stimulus numbers. Proceeding down each column, the participants performed the same arithmetic task (e.g., adding 6 and reporting the sum verbally) with respect to each stimulus number in the column, or they alternated between two different tasks (e.g., adding 6 to the first stimulus number and reporting the sum verbally, subtracting 3 from the second stimulus number and reporting the difference, adding 6 to the third stimulus number and reporting the sum, etc.). The complexity of the required arithmetic operations was either relatively low (e.g., adding 6 and subtracting 3) or high (e.g., adding 17 and subtracting 13). Mean times taken to complete the columns of stimuli were measured as a function of operation complexity and task alternation versus repetition. High-complexity operations took longer on average. Task switching also increased the mean completion times. These two effects interacted reliably; the difference between completion times for task alternation and repetition was greater when the tasks required high-complexity operations.

According to the logic of Sternberg's (1969) additive-factor method, this interaction suggests that operation complexity and task switching influence at least one stage of processing in common. The affected stage may involve some type of executive control process. For example, it might serve to activate the rules used in performing each successive task.

Spector and Biederman's (1976) Study

Extending Jersild's (1927) research, a further study with various versions of the successive-tasks procedure was conducted by Spector and Biederman (1976). It revealed that the sizes of switching-time costs depend on visual cues about what task should be performed next. This dependence suggests that there is an executive control process through which such cues are used along with other stored information to identify and prepare for impending tasks.

In one experiment, Spector and Biederman gave participants columns of two-digit stimulus numbers. For each column, the participants added 3 to every stimulus number and reported the sum verbally, subtracted 3 from every stimulus number and reported the difference, or alternated between adding and subtracting 3. No visible cues were presented to indicate which arithmetic operation should be performed next; instead, the relevant operations had to be recalled from memory. Under these conditions, task alternation took substantially more time than task repetition, as Jersild (1927) found.

In another experiment, Spector and Biederman modified their procedure, appending explicit visual cues (e.g., "+3" or "−3") to the stimuli that indicated which arithmetic operations should be performed. Task alternation still took extra time, but the switching-time cost was markedly lower than when participants were not explicitly cued about the next required operation. The reduction in the switching-time cost could stem from a contribution of task cuing to executive mental control. For example, there may be a control process that identifies what task should be performed next. This process may be facilitated by relevant external information, which helps forego time-consuming memory retrieval.

Yet not all investigators would attribute Spector and Biederman's (1976) or Jersild's (1927) results to anticipatory components of executive mental control. Instead, Allport et al. (1994) hypothesized that time costs of task switching stem from *task-set inertia* (TSI), a type of proactive interference between conflicting S-R mappings for successive tasks. Support for this hypothesis was provided by a study that Allport et al. conducted with the successive-tasks procedure.

Allport et al.'s (1994) Study

The study by Allport et al. (1994) yielded several sets of results (see Table 1). Some of them have been claimed to show that significant TSI occurs and that anticipatory executive control processes play little if any role in task switching. Nevertheless, other results of Allport et al. appear more consistent with executive mental control than with TSI. What led to this ambiguous state of affairs is discussed next in more detail.

Evidence against executive control processes. Some putative evidence against the importance of executive control processes for task switching is that switching-time costs may not depend on the scope of the switches (see Table 1, Result A1). Allport et al. (1994, Experiment 1) found this by presenting visual stimulus displays that contained multiple copies of a particular printed digit. For these displays, participants performed four alternative tasks with different S-R mappings defined by which stimulus attributes and response criteria were relevant. The tasks involved (a) saying whether the magnitude (i.e., absolute value) of the displayed digit was "odd" or "even," (b) saying whether the digit's magnitude was "more" or "less" than 5, (c) saying whether the numerosity of the digit's copies was "odd" or "even," or (d) saying whether this numerosity was "more" or "less" than 5. When participants alternated between two tasks that differed only in their relevant stimulus attributes (i.e., magnitude vs. numerosity), the mean RT was about 1,100 ms longer than for repetitive task performance. Approximately the same 1,100-ms switching-time cost occurred for alternations between two tasks that differed only in their response criteria (i.e., odd–even vs. more–less). Moreover, the switching-time cost was approximately the same for alternations between two tasks that differed in both their relevant stimulus attributes and response criteria. Widening the scope of the switches did not increase their time cost significantly.

According to Allport et al. (1994), this invariance suggests that a "unitary central executive" does not mediate task switching. Their theoretical interpretation assumed that executive control processes have limited capacity and that more mental workload is

[2] Jersild's (1927) study involved a precursor rather than prototypical version of the successive-tasks procedure in that his stimulus displays did not preclude concurrent encoding of multiple stimuli for which responses had to be made successively.

Table 1
Results From Allport et al. (1994) Concerning Executive Control Processes Versus Task-Set Inertia (TSI) in Task Switching

A. Evidence against executive control processes
1. Switching-time costs may not depend on the scope of the switches.
2. Switching-time costs may be present after very long response–stimulus intervals.
3. Switching-time costs may be absent after very short response–stimulus intervals.
B. Evidence for TSI
1. Switching-time costs may be small when stimulus–response mappings are dissimilar.
2. Switching-time costs are increased by prior experience with currently irrelevant tasks.
C. Evidence against TSI
1. Switching-time costs may be absent when TSI should be present.
2. Switching-time costs may be present when TSI should be absent.
D. Evidence for executive control processes
1. Switching-time costs may be unaffected by within-task difficulty.

imposed on these processes by switching between tasks whose relevant stimulus attributes and response criteria both differ. If so, then switching-time costs should increase with greater workload. However, such an increase did not occur empirically, which led Allport et al. to conclude that task switching involves time-consuming processes other than executive mental control per se.

More putative evidence against the importance of executive control processes for task switching is that persistent switching-time costs may occur as RSIs increase (see Table 1, Result A2). Allport et al. (1994, Experiment 5) found this by having participants perform four more tasks that involved alternative S-R mappings: vocally naming the colors of fonts in which different color words (e.g., the word *red* with blue ink) were printed (the *standard Stroop task;* MacLeod, 1991; Stroop, 1935), naming the colors of fonts in which rows of Xs were printed (the *Stroop control task*), reading color words that were printed in fonts of different colors (the *reverse Stroop task*), and reading color words that were printed in black font (the *reverse Stroop control task*). For the reverse Stroop task, mean RTs were longer when it was performed in alternation with the standard Stroop task than when it was performed repetitively. This difference changed relatively little with the RSI. After RSIs of 20 and 1,100 ms, mean switching-time costs were about 180 ms and 135 ms, respectively.

From these results, Allport et al. (1994) again inferred that task switching entails little, if any, anticipatory executive mental control. They reasoned as follows. Suppose that executive control processes do mediate task switching and that these processes commence at the start of the RSI. Also, suppose that switching-time costs stem from the duration of these processes. Then after relatively short RSIs, some switching-time cost should occur. In contrast, after longer RSIs, there should be no switching-time cost; the executive control processes should finish before the next task's stimulus is presented, which would preclude them from contributing to RTs for the next task. However, for the reverse Stroop task, this expected pattern of results failed to occur (Allport et al., 1994, Experiment 5), casting doubt on whether executive control processes mediated switching to this task.

Also relevant is a further result from performance of the reverse Stroop and standard Stroop tasks (Allport et al., 1994, Experiment 5). Mean RTs for the standard Stroop task manifested almost no switching-time costs even when the RSI was very short (see Table 1, Result A3). Allport et al. (1994, Experiment 4) found other cases of null switching-time costs as well. This seems hard to reconcile with executive control processes that play a prominent anticipatory role in task switching. If task switching involves these processes, then after short RSIs, should they not generally yield substantial switching-time costs?

Evidence for task-set inertia. To explain why task switching sometimes but not always produces substantial switching-time costs, Allport et al. (1994) proposed the TSI hypothesis. It is based on two assumptions: (a) Performance of a prior task requires imposing a particular "task set" that increases the primacy of the task's S-R mapping and may also suppress other competing S-R mappings, and (b) the prior task's S-R mapping remains partially active even after long RSIs, potentially interfering with selection of responses for other subsequent tasks. According to Allport et al., this proactive interference is higher when the stimuli and responses for prior and subsequent tasks are similar and when a prior task involves a less dominant S-R mapping than does a subsequent task.[3]

Some putative evidence for the TSI hypothesis is that switching-time costs may be very small when participants alternate between two tasks whose S-R mappings are dissimilar (see Table 1, Result B1). For example, such a result occurred in the first phase of another experiment by Allport et al. (1994, Experiment 4). Here a new group of participants performed only two tasks at the outset: reverse Stroop and digit-magnitude judgment. The stimuli, responses, and S-R mapping for each task differed from those of the other task. After participants had completed a few blocks of

[3] It may seem counterintuitive that higher proactive interference would occur when a prior task involves a less dominant S-R mapping. Nevertheless, Allport et al. (1994, p. 442) have proposed a rationale for how this could happen. According to them, the previously imposed task set must be especially strong when the prior task is less dominant, increasing the degree to which the prior task's S-R mapping remains active and continues to interfere with the use of competing mappings for other subsequent tasks.

practice with the tasks, switching-time costs approached zero; they were much lower than when previous participants had alternated between the reverse Stroop and standard Stroop tasks (cf. Allport et al., 1994, Experiment 1). Under the TSI hypothesis, this is what should have occurred given that proactive interference presumably influenced task switching in the former but not the latter experiment.

A second piece of putative evidence for the TSI hypothesis is that prior experience with related but currently irrelevant tasks may increase switching-time costs considerably (see Table 1, Result B2). For example, this result occurred in a further phase of the experiment described earlier (Allport et al., 1994, Experiment 4). Here participants again performed the same two (reverse Stroop and magnitude judgment) tasks. In the interim, however, they performed two new tasks (Stroop color naming and digit-numerosity judgment) that were related to the preceding ones. After this additional experience, a large increase occurred in the switching-time costs for the reverse Stroop and magnitude-judgment tasks. Such an outcome follows naturally from the TSI hypothesis; residual proactive interference from the S-R mappings of the intervening new tasks may have impeded subsequent performance of the reverse Stroop and magnitude-judgment tasks on alternating-task blocks.

Evidence against task-set inertia. Nevertheless, there is also considerable evidence against the TSI hypothesis. Under some conditions, almost no switching-time cost may occur when TSI should be present (see Table 1, Result C1). For example, let us again consider what happened when Allport et al. (1994, Experiment 5) had participants alternate between the reverse and standard Stroop tasks. Here mean word reading RTs were significantly longer than those obtained when participants alternated between the corresponding control tasks. This result suggests that in alternating-task blocks, having to name ink colors for the standard Stroop task caused interference with subsequent word reading for the reverse Stroop task. Thus, under the TSI hypothesis, participants should have suppressed color naming and imposed a word reading task set for the reverse Stroop task in alternating-task blocks. In turn, such regulation should have caused a significant time cost for switching back to the standard Stroop task, which involves color naming rather than word reading. However, contrary to this prediction, the mean switching-time cost for the standard Stroop task was virtually nil.

Further evidence against the TSI hypothesis is that significant switching-time costs may occur when TSI should be absent (see Table 1, Result C2). In particular, this occurred during performance of the control tasks for which participants named colored patches and read color words printed with black ink (Allport et al., 1994, Experiment 5). There, task alternation took longer than task repetition, even though the required S-R mappings should not have interfered with each other (i.e., their respective stimulus sets had no shared perceptual features). Such results have been obtained as well under other conditions in which TSI was presumably absent (e.g., Allport et al., 1994, Experiment 3). It therefore appears that some source other than TSI contributes to the time cost of task switching.

What might this other source be? Of course, one possibility is executive mental control. Even if the S-R mappings for two different tasks are dissimilar, supervisory shifts of task set may be required to alternate between them.

Evidence for executive control processes. Additional evidence for the existence and separability of executive control processes is that switching-time costs, although substantial in size, may be unaffected by manipulations of within-task difficulty (Table 1, Result D1). For example, the Stroop and numerosity-judgment tasks of Allport et al. (1994, Experiment 3) yielded considerably longer mean RTs than did the reverse Stroop and magnitude-judgment tasks. However, these RT differences were about the same on alternating-task and repetitive-task blocks; the difficulty of the individual tasks did not affect mean switching-time costs significantly. Allport et al. (1994, Experiments 1–3) also reported several other cases of switching-time costs that were unaffected by task difficulty. Given the logic of Sternberg's (1969) additive-factor method, such data suggest that task switching and task difficulty may influence temporally separate, functionally independent stages of processing. Perhaps executive control processes mediate the effects of task switching, whereas other subordinate processes (e.g., stimulus identification, response selection, and movement production) mediate the effects of task difficulty.

Still, to maintain the latter theoretical interpretation, Allport et al.'s (1994) other results must be reconciled with it. For example, why did their long RSIs not eliminate the time cost of task switching? In answer, a study by Rogers and Monsell (1995) is relevant.

Rogers and Monsell's (1995) Study

Rogers and Monsell (1995) used a version of the successive-tasks procedure called the *alternating-runs paradigm.* During each trial block, runs of two or more successive trials for one task alternated with runs of two or more trials for another task. One task involved pressing keys to indicate whether printed digits were odd or even; the other task involved pressing keys to indicate whether printed letters were consonants or vowels. The stimulus display on each trial contained two characters, one relevant and the other irrelevant for the current task. Some of the irrelevant characters were either congruent or incongruent with impending responses; they came from the stimulus ensemble of the noncurrent task and corresponded respectively to keypresses that would be correct or incorrect for the current task. Other irrelevant characters were neutral (i.e., they did not come from the stimulus ensemble of either task). The spatial location of the stimulus display cued participants about which task should be performed next. RTs were measured as a function of the RSI and other stimulus factors. From these measurements, several instructive findings about the nature of executive control processes emerged.

Irrelevant-character effects. Incongruent irrelevant characters induced the largest switching-time costs (Rogers & Monsell, 1995, Experiment 1). This is consistent with the TSI hypothesis (Allport et al., 1994). One would expect proactive interference from a previously applicable S-R mapping to be highest for such stimulus displays, thereby slowing responses especially on trials that require task switching.

However, the TSI hypothesis cannot explain other results of Rogers and Monsell (1995) so well. For example, a substantial switching-time cost also occurred in the context of neutral irrele-

vant characters, even though they presumably induced no proactive interference with the current task. What might the source of this particular cost be? A possible answer is that executive control processes are needed to switch between tasks regardless of which irrelevant characters appear in a stimulus display.

Response–stimulus interval effects. The latter possibility may be evaluated further from patterns of RSI effects. Under some conditions, Rogers and Monsell (1995, Experiment 2)—like Allport et al. (1994, Experiment 5)—found that switching-time costs did not vanish as RSIs increased. This occurred when the lengths of the RSIs varied within trial blocks. Nevertheless, when the RSIs all had the same length within a trial block, but their lengths varied between blocks, switching-time costs were substantially lower after longer RSIs (Rogers & Monsell, 1995, Experiment 3). The blocked-RSI effect was approximately additive with the irrelevant-character effect on switching-time costs.

On the basis of these results, three conclusions can be reached (Rogers & Monsell, 1995). First, RSI and irrelevant-character effects on switching-time costs may occur during distinct substages of executive mental control. Second, if the RSI is predictable, then it may be used for completing some of the operations needed to switch between tasks. Third, if an RSI is unpredictable, then these operations may be postponed until after the next task's stimulus has appeared.

Still, like Allport et al. (1994), Rogers and Monsell never found that the time costs of task switching entirely vanished after long RSIs. Even during trial blocks with constant 1,200-ms RSIs, which provided ample opportunity for executive control processes to complete their anticipatory operations, there were reliable switching-time costs. This persistence might be attributed either to residual TSI or to executive mental control that is postponed until after the RSI has ended.

Task serial-position effect. To test these possibilities further, Rogers and Monsell (1995, Experiment 6) used trial blocks with alternating runs of four trials per task. They reasoned that if the TSI hypothesis were correct, then proactive interference from a prior task should decay gradually, slowing responses not only in the first but also in the second and perhaps even third serial positions of each four-trial run. However, no evidence of gradually decaying proactive interference was obtained. Mean RTs in the second, third, and fourth serial positions of the four-trial runs were virtually identical to each other and all reliably shorter than the mean RT in the first serial position. Thus, it appeared as though, on each four-trial run, the switching-time cost may have stemmed from executive control processes that completed their operations before the first trial of the run ended.

Theoretical interpretation. On the basis of their results, Rogers and Monsell (1995) proposed a model of task switching with two distinct types of executive control: *endogenous* and *exogenous.* According to this model, endogenous control takes place in a flexible top-down manner, executing anticipatory operations for impending tasks during predictable RSIs. These operations decrease switching-time costs as the RSIs increase, accounting for blocked-RSI effects. However, they leave the system in a partially unprepared state. Exogenous control, which completes final preparations for the next task, is triggered by the onset of the next task's stimulus. The occurrence of the exogenous control process after stimulus onset could yield irrelevant-character effects. The temporal separation of exogenous and endogenous control processes may also account for why irrelevant-character and RSI effects on switching-time costs are approximately additive (Lauber, 1995). Although Rogers and Monsell did not specify exactly what these processes do, more conclusions about them may be reached through the new experiments that we report in this article.

Other Relevant Studies

Rogers and Monsell's (1995) theoretical ideas have also been reinforced by some other studies. For example, Meiran (1996) gave visual precues to participants during a modified version of the successive-tasks procedure, informing them explicitly about what their next task would be. On trials that required task switches, the precues reduced switching-time costs more when the RSIs were long than when they were short. This supports the assumption of endogenous executive control. More such support has been reported by other investigators (e.g., Biederman, 1973; LaBerge, Petersen, & Norden, 1977; Logan & Zbrodoff, 1982; Sudevan & Taylor, 1987).

In addition, complementary evidence of exogenous executive control has been reported by Gopher, Armony, and Greenshpan (2000). Again using visual precues, they had participants make occasional unpredictable switches between tasks. Significant switching-time costs occurred even though the precues were presented at the start of long (1,200-ms) RSIs. However, these costs did not extend beyond the particular trials on which the switches took place; there appeared to be no residue of gradually decaying TSI. This result, reminiscent of Rogers and Monsell's (1995, Experiment 6) all-or-none task serial-position effect, is what would be expected if an exogenous control process completes its task-set shifting immediately and fully after the onset of the next task's stimulus.

Interim Summary

Although the difficulty of task switching may be partly attributable to sources (e.g., TSI) other than executive mental control per se, our literature review suggests that both endogenous and exogenous control processes probably help supervise task switching and contribute significantly to observed switching-time costs. These contributions can account for patterns of effects by factors such as task cuing, operation complexity, and RSI. Thus, further efforts to formulate and test detailed models of executive control for task switching are presumably warranted.

A Stage Model of Executive Control for Task Switching

Given the preceding considerations, the purpose of the present article is to formulate and test a model of executive control that accounts more fully for task cuing, operation complexity, RSI, and other related factor effects on the time costs of task switching. In what follows next, our model's assumptions are outlined. After this, we discuss how they can explain various results in the task-switching literature. Then four new experiments are reported to evaluate some additional predictions of the model.

A schematic diagram of the model appears in Figure 1. According to the model, performance during the successive-tasks proce-

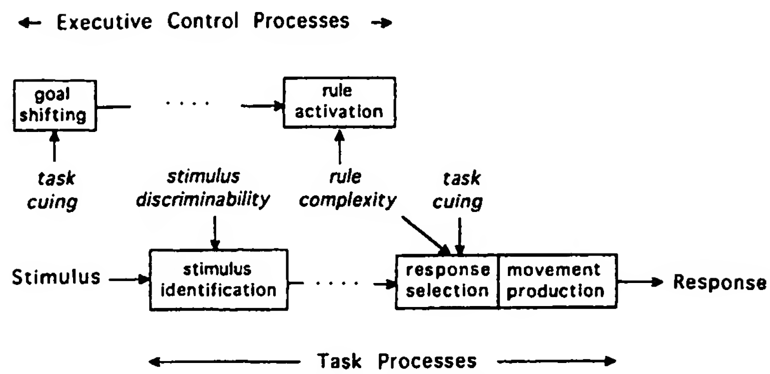

Figure 1. A stage model of task switching that has distinct executive control and task processes. Various factors (viz., task cuing, rule complexity, and stimulus discriminability) may influence the durations of these processes, thereby affecting mean reaction times and switching-time costs additively or interactively.

dure entails two complementary sets of stages: executive control processes and task processes (for more discussion about each type of process, see Lauber, 1995, and Kieras et al., 2000).

Task Processes

We assume that task processes are used for performing individual perceptual–motor and cognitive tasks under both single-task and multiple-task conditions. In our model, these processes include three principal stages, *stimulus identification, response selection,* and *movement production,* which operate on the basis of information in declarative and procedural working memory (cf. Donders, 1868/1969; Meyer & Kieras, 1997a, 1997b; Sanders, 1980; Sternberg, 1969). The stimulus-identification stage encodes perceptual features of stimuli and places them in declarative working memory for access during the response-selection stage. Through algorithms in procedural working memory, the response-selection stage converts the stimulus codes to abstract response codes. The movement-production stage converts the response codes to motor commands that generate overt physical action. Component operations in each stage are assumed to be tailored to the tasks' particular sensory modalities, response modalities, and S-R mappings.

Regarding the response-selection stage, we further assume that it uses production rules in procedural working memory, which specify actions to be executed whenever prerequisite conditions match the current contents of declarative working memory. For example, suppose that a task requires pressing finger keys in response to stimulus colors. Then a production rule for response selection might have the following form:

IF ((GOAL IS TO DO COLOR-DISCRIMINATION TASK) AND (STIMULUS COLOR IS RED))
THEN (PRESS RIGHT INDEX-FINGER KEY).

The numerosity and complexity of such rules depend on the task's S-R mapping, thereby affecting the duration of the response-selection stage (Meyer & Kieras, 1997a, 1997b).

When the same task is performed repetitively, response selection in the model starts immediately after stimulus identification on each trial. However, we assume that if a switch occurs from one task to another, there is a pause between the end of stimulus identification and the beginning of response selection for the current task (see Figure 1). This pause is used by an executive control process whose operations enable the subsequent response-selection stage to proceed correctly.

Executive Control Processes

To enable task switching, the model's executive control processes include two distinct stages, *goal shifting* and *rule activation,* which are accomplished through executive production rules. Together, goal shifting and rule activation respectively ensure that the contents of declarative and procedural working memory are appropriately configured for the task at hand, consistent with proposals of some previous theorists (e.g., Duncan, 1986; Kimberg & Farah, 1993; Logan, 1985; Meyer & Kieras, 1997a, 1997b; Rogers & Monsell, 1995).[4]

Goal shifting. The goal-shifting stage keeps track of current and future tasks, inserting and deleting their goals in declarative working memory as needed. Specific goal items in working memory let other components of the system "know" what the current task is. For example, in switching from a shape-discrimination to a color-discrimination task, goal shifting might involve updating

[4] Of course, there are probably other important executive control processes. For example, one may involve attention refocusing, which tunes particular perceptual mechanisms to be especially sensitive in discriminating relevant stimulus features for the current task (Keele & Rafal, 2000; Meyer et al., 1997, 1998). Also, executive control may involve response monitoring and error detection to help ensure that overt movements are produced correctly despite underlying system noise (Gehring, Goss, Coles, Meyer, & Donchin, 1993; Kornblum, Hasbroucq, & Osman, 1990; Norman & Shallice, 1986). This supervision could be especially important during the successive-tasks procedure, which entails conflicting S-R mappings. For now, however, we focus on the goal-shifting and rule-activation stages of task switching.

the contents of working memory through the following production rule:

> IF ((GOAL IS TO DO SHAPE-DISCRIMINATION TASK) AND (SHAPE-DISCRIMINATION TASK IS DONE) AND (NEXT TASK IS COLOR DISCRIMINATION))
> THEN (((DELETE (GOAL IS TO DO SHAPE-DISCRIMINATION TASK)) AND (INSERT (GOAL IS TO DO COLOR-DISCRIMINATION TASK))).

By the application of such rules, various bits of information for initiating, executing, and terminating individual tasks can be maintained.

Furthermore, we assume that the time at which goal shifting takes place relative to concomitant task processes is flexible. Under some conditions, goal shifting may occur before stimulus identification starts for the next task (see Figure 1). For example, this might happen if the RSI is long and prior information is available about what the next task will be. Then the goal-shifting stage would be an endogenous control process of the sort Rogers and Monsell (1995) have proposed. However, our model also allows goal shifting to occur after the next task's stimulus has been identified. Such delayed goal shifting might occur if the RSI is short or the task stimulus is expected to provide an explicit cue about what task must be performed next. Then goal shifting would be an exogenous control process.

Rule activation. In the model, rule activation is another executive control process for task switching. Because of reasons explained subsequently, we assume that under at least some conditions, this stage is triggered exogenously and takes place during a pause between the end of stimulus identification and the beginning of response selection for the current task, after goal shifting has finished (see Figure 1). Two complementary functions are served by rule activation: enabling the rules for selecting the current task's response and disabling the rules for selecting the prior task's response. After these functions have been completed, the current task's response-selection stage can proceed.

How is rule activation accomplished? One possibility is that this stage involves "loading" the next task's rules into procedural working memory, just as a computer operating system reads new application programs from disk to core memory, overwriting old programs in preparation for executing the new ones. Such operations might be initiated by the following executive production rule:

> IF ((GOAL IS TO DO COLOR-DISCRIMINATION TASK) AND (STIMULUS COLOR HAS BEEN IDENTIFIED))
> THEN (((LOAD (COLOR-DISCRIMINATION TASK RULES)) AND (INSERT (WAIT FOR COMPLETION OF LOADING))).

On the basis of the loading-of-rules metaphor, it seems plausible that the numerosity and complexity of a task's production rules could influence the duration of the rule-activation stage.

Another complementary possibility is that this stage involves temporarily raising the activation levels of the current task's production rules in procedural long-term memory (cf. J. R. Anderson, 1983, 1993). During such a process, the activation levels of the previous task's rules might be allowed to drop back toward baseline, or an intentional operation to suppress them might occur (cf. Goschke, 2000; Mayr & Keele, 2000). If either of these possibilities holds, then procedural working memory would constitute the part of procedural long-term memory that is currently activated. This could help account for putative effects of TSI on task switching (cf. Allport et al., 1994; Allport & Wylie, 2000).

Theoretical rationale. Of course, our assumptions about rule activation lead to other questions. Why is this stage necessary? Why are the production rules for multiple tasks not kept simultaneously enabled in procedural working memory during the successive-tasks procedure? Why might rule activation be an exogenous rather than endogenous control process?

One conceivable answer to some of these questions is that procedural working memory has only enough capacity for a single task's rules. However, this seems implausible. Results from previous studies with the PRP simultaneous-tasks procedure suggest that, under at least some conditions, sets of rules for two distinct tasks can be held in procedural working memory and used concurrently during multiple-task performance (Meyer & Kieras, 1997a, 1997b). Thus, there must be some other rationale for our proposed rule-activation stage.

In particular, perhaps rule activation is needed because of the successive-tasks procedure's special characteristics. As mentioned before, this procedure typically involves tasks that have the same stimuli but different S-R mappings. Under such conditions, it may be suboptimal to keep the production rules for all of the tasks enabled in procedural working memory. Doing so could create conflict, disruption, and errors during the response-selection stage for one task or another, because alternative rules whose stimulus conditions are satisfied simultaneously would yield inappropriate—and even mutually exclusive—actions (cf. Cohen, Dunbar, & McClelland, 1990; Kornblum, Hasbroucq, & Osman, 1990; MacLeod, 1991). One solution to these problems would be to enable the rules for only one task at a time, as our rule-activation stage does.

Given these considerations, proactive interference and TSI might play a significant role during rule activation (cf. Allport et al., 1994; Allport & Wylie, 2000). Suppose that some of the features of the next task's stimulus match those of the conditions in a production rule for the preceding task. Also, suppose that the matching features enter declarative working memory before rule activation has finished for the next task. Then the occurrence of such partial matches could make it more difficult to disable the preceding task's rules, thereby prolonging the rule-activation stage (Mayr & Keele, 2000).

These considerations could also justify having the rule-activation stage be an exogenous (stimulus-triggered) control process. Perhaps attending to relevant features of the next task's stimulus helps determine which production rules should be enabled for dealing with it. If so, then on each trial that involves task switching, rule activation might benefit from waiting until the stimulus for the next task has been identified.

Ancillary Technical Assumptions

To derive explanations and predictions from the present model, we make some further technical assumptions that are commonly

associated with discrete stage models of human information processing (Sanders, 1980; Sternberg, 1969).

Strict successiveness. The model's component stages, including both executive control and task processes, are strictly successive. Each stage starts only after its predecessors have finished.

Summation of stage durations. Theoretical RTs are sums of component stage durations. On trials without task switching, the summed durations of stimulus identification, response selection, and movement production constitute the RTs. On trials with task switching, these summed durations in combination with those of goal shifting and rule activation constitute the RTs.

Selective influence of factors. Some factors may selectively influence the component stage durations of different executive control and task processes. However, other factors may influence multiple stages, and some stages may be influenced by multiple factors.

Constant output quality. The quality of the outputs produced by a component stage is constant regardless of the factor effects on its duration.

Additivity and interaction of factor effects. Factors that selectively influence the durations of different component stages have additive effects on mean RTs. In contrast, the effects of factors that influence the same stage may interact.[5]

Justification of Assumptions

Although discrete stage models have enjoyed considerable popularity (Donders, 1868/1969; Luce, 1986; Meyer, Osman, Irwin, & Yantis, 1988; Miller, 1988; Pachella, 1974; Sanders, 1980; Sternberg, 1969; Townsend & Ashby, 1983), their relevance is conceivably limited. For example, McClelland (1979) argued that human cognition and action are typically mediated by a cascade of contingent–concurrent operations whose outputs consist of continuous, gradually increasing activation. Similar arguments might be made by theorists who favor connectionist-network models (e.g., Cohen et al., 1990). If they were empirically correct, then our model's assumptions would not strictly hold. Nor would Sternberg's (1969) additive-factor method be entirely applicable here.

Nevertheless, we have strong grounds for initially adopting a discrete stage model and the additive-factor method. In other domains of cognitive psychology (e.g., studies of visual word recognition), such models have been especially useful even when later replaced by alternative theoretical frameworks (Meyer, Schvaneveldt, & Ruddy, 1975). Their simplicity and rigor provide powerful heuristics for conceptualizing basic processes of human performance. Furthermore, discrete stage models account well for a wide range of RT data (Roberts & Sternberg, 1993; Sternberg, 1969, 1998). Thus, it seems likely that they may yield important insights about executive control processes and task switching as well.

Explanation of Past Findings About Task Switching

Our stage model of task switching explains a variety of findings from past studies with the successive-tasks procedure. Many reported differences in switching-time costs may be attributed to factor effects on either goal shifting or rule activation. Perhaps such effects have sometimes been additive because of their disparate loci in the hypothesized sequence of processing stages.

Effects on goal shifting. For example, one factor that probably affects the goal-shifting stage is task cuing. As mentioned earlier, Spector and Biederman (1976, Experiments 3 and 4) found smaller switching-time costs when alternative types of arithmetic problems were accompanied by corresponding operation signs. This decrease may have occurred because the operation signs cued participants with useful information about which task goal should be placed in declarative working memory next, thereby shortening the number of steps taken to complete goal shifting.

A second factor whose effect probably occurs in goal shifting is the length of the RSI. Insofar as the RSI is relatively long or short, it would allow more or less of this stage to be completed before the onset of the stimulus for the next task. Consequently, goal shifting's contribution to RTs could be less when the RSI is long, reducing switching-time costs as Rogers and Monsell (1995, Experiment 3) found with blocked RSIs. This explanation would also account for results of Meiran (1996), who found that task cues reduced switching-time costs more after longer RSIs.

However, the model does not imply that, after long RSIs, switching-time costs should necessarily vanish. On the contrary, suppose that the length of the RSI varies randomly across trials. Then for each trial that requires a task switch, the goal-shifting stage may be postponed until the next task's stimulus is identified. Such optional postponement would preclude RSI effects on switching-time costs, as Rogers and Monsell (1995, Experiment 2) found with mixed RSIs. Also, for each trial that requires a task switch, rule activation may occur after the RSI has ended. This stage could therefore yield residual switching-time costs even after long RSIs, as both Allport et al. (1994) and Rogers and Monsell (1995) found.

Effects on rule activation. The rule-activation stage is also a likely site of operation-complexity effects. Recall that Jersild (1927) found greater switching-time costs when participants alternated been complex rather than simple arithmetic operations. This may have occurred because more production rules are needed to perform complex operations, and it takes longer to activate them, just as larger "number crunching" programs take longer to be loaded in a digital computer's memory.

Irrelevant-character effects probably occur during rule activation as well. For example, recall that Rogers and Monsell (1995) found greater switching-time costs when stimulus displays contained incongruent rather than neutral irrelevant characters. This could occur because incongruent irrelevant characters make it

[5] The additivity and interaction of factor effects discussed here would occur regardless of whether RTs from individual trials are composed of stochastically independent stage durations. Our model does not assume that stochastic independence necessarily holds. For example, across trials, the durations of goal shifting, rule activation, stimulus identification, and response selection might correlate positively or negatively with each other, because of systematic fluctuations in arousal, fatigue, and other subjective factors. Yet mean RTs can still be analyzed to reveal the existence, functional separability, and temporal successiveness of these processing stages; additive-factor effects on mean stage durations may occur even if the stage durations on individual trials are stochastically dependent (Sternberg, 1969).

harder to disable the production rules of prior tasks. Such difficulty would likewise explain why Allport et al. (1994, Experiments 4 and 5) found greater switching-time costs when current stimuli contained perceptual features associated with previous familiar tasks.

Certain characteristics of the rule-activation stage might explain other results of Allport et al. (1994, Experiment 1). As mentioned earlier, they found that under some conditions, mean switching-time costs were about the same regardless of the switches' scope (i.e., the costs did not depend on whether the relevant stimulus features, the response ensemble, or both changed between tasks). Perhaps this occurred because the scope of the switches did not affect the numerosity or complexity of the production rules that had to be enabled in task switching and so did not affect the duration of rule activation either.

In essence, rule activation presumably prepares the processing system so that response selection can proceed rapidly for the current task. After this has been accomplished, there need be no further switching-time cost until a subsequent task switch must take place. The immediate completion of the rule-activation stage accounts for why the time cost that Rogers and Monsell (1995, Experiment 6) observed on the first trial after a task switch did not propagate in a gradually decreasing fashion throughout a run of successive trials with the same task.

Additive factor effects. If the preceding explanations are correct, then our model would account for why certain factor effects on switching-time costs have been essentially additive. For example, recall that Rogers and Monsell (1995, Experiment 3) found such additivity in RSI and irrelevant-character effects. This may have occurred because these two factors respectively affect goal shifting and rule activation, which are successive stages whose durations jointly constitute the time cost of task switching.

Similarly, it is possible to account for some additive factor effects that Allport et al. (1994) found. Recall that their participants took more time respectively on the standard Stroop and digit-numerosity judgment tasks than on the reverse Stroop and digit-magnitude judgment tasks (Allport et al., 1994, Experiment 1). This effect of task difficulty was about the same during repetitive-task and alternating-task trial blocks, even though block type affected mean RTs reliably (i.e., the difficulty and block-type effects on mean RTs were additive). Our model explains such additivity because trial-block type may influence stages of executive control, whereas the difficulty of particular tasks may stem from within-task stages (i.e., task processes) such as stimulus identification and response selection. For analogous reasons, the model is also consistent with other additive factor effects found by Allport et al. (1994, Experiments 2 and 3).

Overview of Experiments

The present article reports four experiments involving various versions of the successive-tasks procedure designed to further test several predictions based on our model. First, we show that as the model predicts, executive control and task processes can be empirically dissociated and affected separately by different factors (Experiment 1). Second, we show that executive control entails at least two component stages, goal shifting and rule activation, whose mean durations depend respectively—and additively—on task cuing and rule complexity (Experiment 2). Third, we show that because of how rule activation works, switching-time costs may be asymmetric in ways related to the familiarity of individual tasks between which participants must switch (Experiments 3 and 4). Taken together, the results of the four experiments strongly support the model's basic assumptions about the nature of executive mental control for task switching.

Methodological Approach

During each experiment, repetitive-task blocks of trials were completed for each of two tasks. There were also alternating-task blocks in which participants switched back and forth between the two tasks at hand. Short RSIs were used throughout the experiments, thereby helping to maximize observable contributions of goal shifting and rule activation to switching-time costs.

To show that these executive control processes are functionally distinct and separable from task processes, we manipulated several factors, including type of task, complexity of task rules, availability of task cues, and discriminability of task stimuli with which participants worked. In Experiments 1 and 4, participants rapidly classified visual patterns of geometric objects with respect to alternative perceptual categorization rules. For the pattern-classification tasks, rule complexity was manipulated by having participants apply unidimensional or bidimensional classification rules. In Experiments 2 and 3, participants solved arithmetic problems with alternative numerical combination rules. For the arithmetic tasks, rule complexity was manipulated by requiring addition and subtraction or multiplication and division operations. Visual task cues (i.e., arithmetic-operation signs) were presented during some trial blocks but not others. Our combined manipulations across the experiments allowed us to check for expected patterns of additive and interactive factor effects on mean RTs and switching-time costs, which provide diagnostic indicators of temporally separate processing stages (Roberts & Sternberg, 1993; Sternberg, 1969, 1998).

Estimation of Switching-Time Costs

Following Allport et al. (1994), we estimated mean switching-time cost (T_S) as follows:

$$T_S = [T_{12} - .5(T_1 + T_2)]/(n - 1), \tag{1}$$

where $n - 1$ is the number of task switches in an alternating-task block of n trials; T_1 and T_2 are the mean completion times for repetitive-task blocks of trials with Tasks 1 and 2, respectively; and T_{12} is the mean completion time for alternating-task blocks. According to our stage model, Equation 1 yields an estimate for the summed mean durations of goal shifting and rule activation on alternating-task trials when the RSI is zero. In addition, the mean RTs on repetitive-task trials yield estimates for the summed mean durations of stimulus identification, response selection, and movement production. Thus, by examining factor effects on both mean RTs and switching-time costs,

we may analyze the contributions of various executive control and task processes to task switching.[6]

Experiment 1

Our choice of tasks for Experiment 1 was inspired by the Wisconsin Card Sorting Test (WCST), a neuropsychological test used to assess executive mental control in patients with brain damage (S. Anderson, Damasio, Jones, & Tranel, 1991; Drewe, 1974; Grafman, Jones, & Salazar, 1990; Grant & Berg, 1948; Heaton, 1981; Milner, 1963). As in the WCST, participants in Experiment 1 perceptually classified and sorted stimulus cards so that the size, shape, shading, or numerosity of geometric objects on them matched those of objects on corresponding target cards. Unlike in the WCST, however, the rules for the present visual pattern-classification tasks were fully explained beforehand, and participants knew when to switch from one task to the next. These modifications allowed us to isolate particular stages of processing that the WCST confounds.

Specifically, in Experiment 1, we tested whether executive control processes exist and can be separated from accompanying basic task processes. For this first step, our manipulations involved two factors: stimulus discriminability and rule complexity. Here we show that stimulus discriminability, a factor especially relevant to stimulus identification (Sanders, 1980; Sternberg, 1969), affects mean RTs reliably on both repetitive-task and alternating-task trial blocks but has no reliable effect on mean switching-time costs. This pattern of effects is predicted by our model, assuming that task processes such as stimulus identification are functionally independent of and complementary to executive control processes. Concomitantly, we show that mean RTs on alternating-task blocks exceed those on repetitive-task blocks, as would be expected if distinct executive control processes contribute to task switching. We also show that rule complexity (viz., unidimensional vs. bidimensional classification), a factor particularly relevant to rule activation and response selection, affects mean RTs somewhat on repetitive-task blocks but more so on alternating-task blocks, yielding a reliable effect on mean switching-time costs. Again, this pattern of effects is what the model predicts, assuming that executive control processes such as rule activation exist over and above task processes such as response selection.

Method

Participants. The participants were 12 undergraduates enrolled in introductory psychology courses at the University of Michigan who received course credit for taking part. Each participant performed individually during a 1-hr session. None committed more than 15 errors during test trial blocks.

Laboratory apparatus and environment. Each session was conducted in a quiet, well-illuminated room that contained a small table and two chairs. The participant sat on one side of the table facing the experimenter on the opposite side. There were three phases per session: The experimenter instructed the participant, allowed him or her to practice the tasks at hand, and then administered a series of test trial blocks. Before each block, four target cards were placed on the table in front of the participant, forming a horizontal row that spanned about 50 cm. During a trial block, the participant sorted a deck of index cards that contained the task stimuli. Sorting required holding the stimulus cards face down, turning them face up one after another, and placing them successively into piles that matched the target cards with respect to prespecified categorization rules. The participant's manual movements for sorting were similar to those for dealing a deck of playing cards. A stopwatch was used to measure the participant's total completion time for each block.

Stimuli. Task stimuli were constructed with 14-cm × 11.5-cm index cards laminated in clear plastic. There were 24 stimulus cards per deck. Mounted on each card was a pattern of identical geometric objects. The pattern was defined by a combination of four perceptual dimensions (shape, numerosity, size, and shading of the objects). Across stimulus cards, the objects had four different shapes (triangle, circle, star, and cross), four levels of numerosity (1, 2, 3, and 4), four sizes (4.8, 3.3, 2.3, and 1.5 cm in height), and four levels of shading (dark, medium, light, and white). By construction, the different values along the shape and along the numerosity dimensions were relatively easy to discriminate; the different values along the size and along the shading dimensions were more difficult to discriminate.

The target cards used during a trial block also had patterns of geometric objects on them. For example, Figure 2 shows four targets (Panel A) and a typical stimulus card (Panel B) from the low-rule-complexity condition. The composition of the target cards, and the rules used for sorting the stimulus cards with respect to them, depended on which rule-complexity condition was involved.

Low-rule-complexity condition. For the low-rule-complexity condition, there were four different tasks that required sorting with respect to either the shape, size, shading, or numerosity of objects on the stimulus cards. The same set of four target cards was used in each case. Cards were constructed such that no two of them matched each other on any dimension. Depending on which task was involved, a particular stimulus card could correctly match any one of the four targets.

Four decks of stimulus cards, with 24 cards per deck, were constructed for performing the four different low-complexity tasks during repetitive-task blocks. In each deck, equal numbers of cards matched the relevant perceptual dimension with respect to each of its four values. Also, each stimulus card had values on three irrelevant perceptual dimensions that matched those of the other targets.

Two additional decks of stimulus cards, again with 24 cards per deck, were constructed for performing the low-complexity tasks during alternating-task blocks. Each deck had 12 cards in odd serial positions for sorting with respect to one perceptual dimension and 12 cards in even serial positions for sorting with respect to another perceptual dimension. The pairs of tasks performed for one deck involved dimensions whose values were highly discriminable (shape and numerosity), whereas the pairs of

[6] To motivate their development of the alternating-runs paradigm, Rogers and Monsell (1995) questioned whether unbiased estimates of executive-process durations can be obtained through comparing performance on temporally separate alternating-task and repetitive-task trial blocks. They suggested that under the latter conditions, which characterized Allport et al.'s (1994) and our procedures, the durations of task processes may change across trial-block types. For example, this could occur because on alternating-task blocks, participants are more aroused and must maintain two S-R mappings simultaneously, whereas only one mapping has to be maintained on repetitive-task blocks. However, experiments conducted by Lauber (1995) suggest that estimates obtained with Equation 1 are probably not biased much, if any, by ancillary differences between performance on repetitive-task and alternating-task blocks. Insofar as we find additive effects of various factors on the durations of distinct executive control and task processes, this further disarms Rogers and Monsell's (1995) concerns; it is difficult to see how such persistent and systematic additivities could occur if their pessimistic scenario prevailed significantly in actual practice.

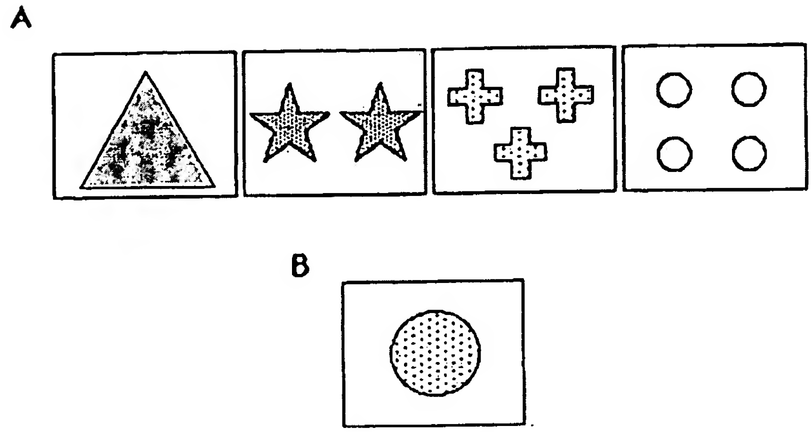

Figure 2. Examples of stimuli in the low-rule-complexity condition of Experiment 1. A: Four target cards used when participants sorted stimulus cards with respect to single perceptual dimensions. The far left target consists of one extra-large dark triangle, the middle left target consists of two large medium-shade stars, the middle right target consists of three medium-size light crosses, and the far right target consists of four small white circles. B: Stimulus card that consists of one large light circle.

tasks performed for the other deck involved dimensions whose values had low discriminability (size and shading).

High-rule-complexity condition. For the high-rule-complexity condition, there were two different tasks that required sorting with respect to either the shape and number or the size and shading of objects on the stimulus cards. Again, four target cards were used in each of these tasks. They were constructed such that a stimulus card matched one target in terms of two relevant perceptual dimensions, two other targets in terms of one relevant dimension, and a fourth target in terms of no relevant dimensions (e.g., see Figure 3).

Because there were four different values on each perceptual dimension, two sets (A and B) of target cards were constructed for the high-complexity tasks. The four targets in Set A respectively displayed one extra-large dark triangle, one medium-size dark cross, three extra-large light triangles, and three medium-size light crosses (see Figure 3, Panel A). The four targets in Set B respectively displayed two large medium-shade stars, two small medium-shade circles, four large white stars, and four small white circles.

Two decks of stimulus cards were constructed for sorting with respect to shape and numerosity on repetitive-task blocks; one deck was used with the target cards of Set A, and the other was used with the target cards of Set B. Similarly, two more decks were constructed for sorting with respect to size and shading on repetitive-task blocks. There were 24 cards per deck, with each target card being matched by 6 of them.

Two additional decks of stimulus cards, again with 24 cards per deck, were constructed for performing the high-complexity tasks during alternating-task blocks. The 12 cards in odd serial positions of a deck were sorted with respect to conjunctions of shape and numerosity; the 12 cards in even serial positions were sorted with respect to conjunctions of size and

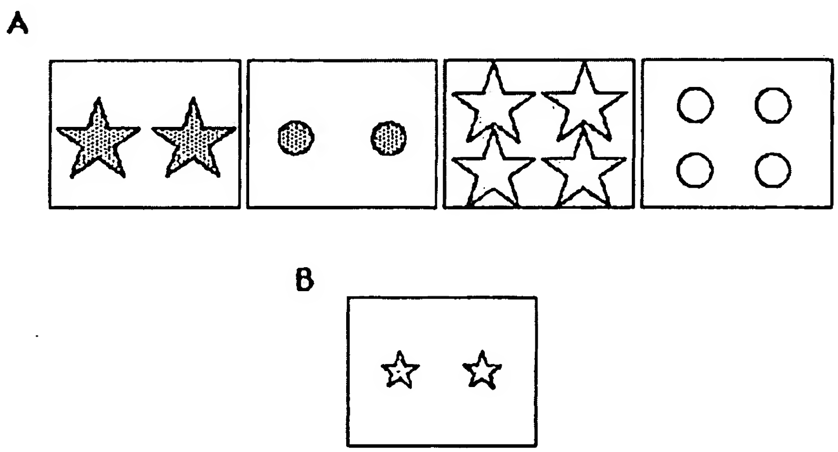

Figure 3. Examples of stimuli in the high-rule-complexity condition of Experiment 1. A: Four target cards used when participants sorted stimulus cards with respect to conjunctions of two perceptual dimensions (shape and numerosity, or size and shading). B: Representative stimulus card.

Table 2
Results of Experiment 1

Trial-block type	Rule complexity	Stimulus discriminability	Relevant dimensions	Mean RT (ms)	Error rate (%)	Switching-time cost (ms)
Repetitive	Low	Low	Size	1,719	5.9	
	Low	Low	Shading	1,424	1.8	
	Low	High	Shape	1,354	1.0	
	Low	High	Numerosity	1,191	0.0	
	High	Low	Size and shading	1,788	0.7	
	High	High	Shape and numerosity	1,715	2.5	
Alternating	Low	Low	Size or shading	2,094	10.4	545
	Low	High	Shape or numerosity	1,830	1.8	582
	High	Low and high	Size and shading, or shape and numerosity	3,080	7.5	1,386

Note. Mean reaction times (RTs) and error rates were calculated by dividing the number of stimulus cards per deck into the total completion times and errors per trial block. Mean switching-time costs were calculated with Equation 1.

shading. Set A provided the target cards for one deck, and Set B provided the target cards for the other. Because the alternating-task blocks for the high-rule-complexity condition necessarily involved all four relevant perceptual dimensions, it was not possible to manipulate stimulus discriminability and trial-block type orthogonally under this condition.

Design. There were 12 blocks of test trials per participant: 4 repetitive-task blocks with the low-complexity tasks (1 for each perceptual dimension), 2 alternating-task blocks with the low-complexity tasks (1 for switching between shape and numerosity and 1 for switching between size and shading), 2 repetitive-task blocks with the high-complexity tasks that involved conjunctions of shape and numerosity (1 block for Set A and 1 for Set B), 2 repetitive-task blocks with the high-complexity tasks that involved conjunctions of size and shading (1 block for Set A and 1 for Set B), and 2 alternating-task blocks with the high-complexity tasks (1 for Set A and 1 for Set B). Stimulus cards were ordered within each deck such that no more than three successive cards matched the same target; duplicate stimulus cards occurred during different halves of a block. All of the target-card sets and stimulus-card decks were used for each participant. The serial orders of the different block types, rule-complexity conditions, and stimulus-card decks were counterbalanced across participants with a nested Latin square design.

Procedure. Brief (5-min) instructions and practice were given at the start of each session. The experimenter showed the participant the target cards and several practice stimulus cards. The different values on each perceptual dimension were described along with the four low-complexity and two high-complexity tasks. The participant practiced these tasks until he or she performed them correctly. Also, to provide more practice in rapid card sorting, four target cards labeled *A, B, C,* and *D* were aligned in front of the participant. With respect to these targets, the participant sorted two decks of practice cards that had various single letters on them.

After practice, the blocks of test trials commenced. At the start of each block, the experimenter informed the participant about what the task (or tasks) would be. The participant was told to sort the stimulus cards as quickly as possible without making many errors. Total completion time for sorting was recorded on each trial block. Timing began when the experimenter said "Start" and ended when the participant put the last card of a deck on the table. The experimenter also recorded any errors that occurred along the way.

Results

Table 2 shows mean RTs, error rates, and switching-time costs as a function of trial-block type, rule complexity, and stimulus discriminability for each combination of visual pattern-classification tasks in Experiment 1. Error rates were low on average (4.3% across conditions) and correlated positively with mean RTs ($r = .69$, $p < .05$), suggesting that systematic RT differences did not stem simply from speed–accuracy trade-offs across conditions. The reliability of factor effects on mean RTs and switching-time costs was evaluated through repeated measures analyses of variance.[7]

Effects of trial-block type. Trial-block type affected mean RTs reliably, $t(11) = 13.8$, $p < .0001$. On average, participants took more time to respond during alternating-task blocks than during repetitive-task blocks of pattern classifications, manifesting substantial switching-time costs ($M = 975$ ms, $SE = 71$ ms), which would be expected if—as our model predicts—executive control processes contribute significantly to task switching.

Effects of rule complexity. Rule complexity affected mean RTs during both repetitive-task and alternating-task blocks (see Figure 4, left panel). On average, high-complexity rules yielded slower pattern-classification responses than did low-complexity rules (mean difference = 724 ms, $SE = 72$ ms), $t(11) = 9.99$, $p < .0001$. The magnitude of this effect on mean RTs during repetitive-task blocks (viz., 330 ms) presumably manifests how much longer judgment and response selection took with high-complexity (i.e., bidimensional classification) rules than with low-complexity (i.e., unidimensional classification) rules. For now, we assume that an equivalent lengthening of these task processes occurred during alternating-task blocks, which is supported by subsequent evidence that task and executive control processes may be functionally dissociable from each other.

Nevertheless, during alternating-task blocks, rule complexity had a greater total effect. As a result, switching-time costs were

[7] The present analyses concern main effects and interactions that have one degree of freedom. Their reliability was quantified originally in terms of F values. However, for purposes of exposition, we have transformed them to t values. This allowed us to test unidirectional (one-tailed) as well as bidirectional (two-tailed) hypotheses involving a priori predictions about the signs (positive or negative) that particular differences between mean RTs and switching-time costs should have. Through squaring of the reported t values, they may be transformed back to F values whose numerators have one degree of freedom and whose denominators have the same degrees of freedom as the t values (Hays, 1963). Several ancillary factors whose levels were counterbalanced in our design did not have reliable effects, so the reported analyses averaged our data across them.

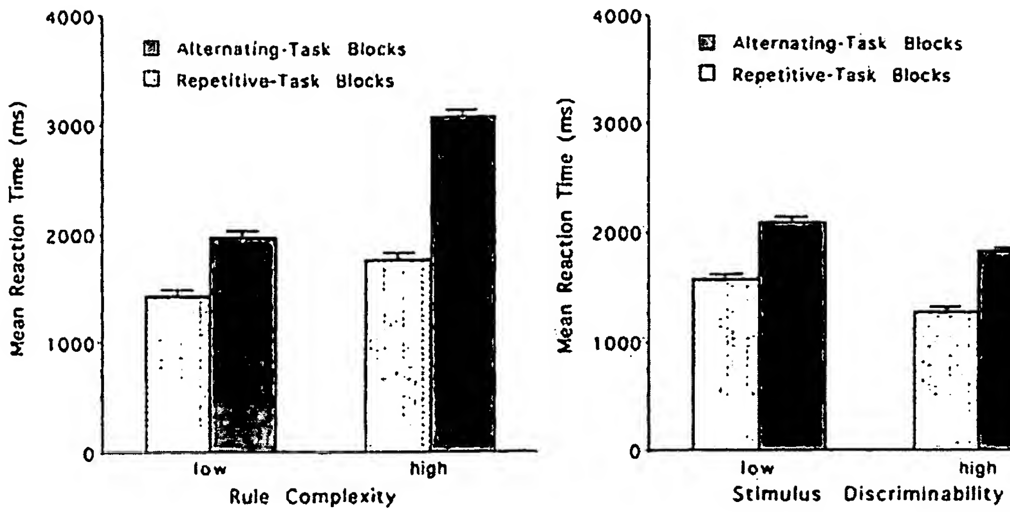

Figure 4. Results of Experiment 1. Left: Mean reaction times (RTs) as a function of rule complexity and trial-block type. Standard errors (lines extending above vertical bars) are based on the interaction among block type, rule complexity, and participants. Right: Mean RTs as a function of stimulus discriminability and trial-block type in the low-rule-complexity condition. Standard errors are based on the interaction among block type, stimulus discriminability, and participants.

reliably greater for high-complexity pattern classifications than for low-complexity pattern classifications (mean difference = 823 ms, SE = 128 ms), $t(11)$ = 6.43, p < .0001. This supports the assumption of a rule-activation stage in executive control that enables task processes such as response selection. As predicted by our model, rule activation apparently makes a significant contribution to task switching, and its duration depends on the complexity of the rules for the next task that has to be performed.

Effects of stimulus discriminability. Mean RTs were also affected reliably by stimulus discriminability, whose effects presumably occurred during the task process of stimulus identification. For low-complexity pattern classifications, responses required less time on average when the relevant perceptual dimension had highly discriminable values than when its values were less discriminable (mean difference = 282 ms, SE = 43 ms), $t(11)$ = 6.50, p < .0001. However, this effect was about the same during both repetitive-task and alternating-task blocks (Figure 4, right panel). Switching-time costs did not change much as a function of the discriminability factor (mean difference = −37 ms, SE = 82 ms), $t(11)$ = −0.45, p > .5. This null result supports our model's prediction that task and executive control processes can be influenced selectively by different experimental factors.[8]

Discussion

From the results of Experiment 1, we conclude that our stage model of executive control in task switching merits further consideration. Conforming to the model, mean RTs were longer during alternating-task blocks than during repetitive-task blocks of the visual pattern-classification tasks. The effects of rule complexity interacted with those of trial-block type, as should happen if there is a rule-activation stage whose duration constitutes one component of switching-time costs. In contrast, stimulus discriminability also reliably affected mean RTs but not switching-time costs, confirming the predicted separability of executive control and task processes such as stimulus identification. Given these findings, it remains to be established that executive control also includes a distinct stage of goal shifting, which has been assumed in our model (see Figure 1).

Conceivably, the TSI hypothesis of Allport et al. (1994) could also account at least qualitatively for the results of Experiment 1. Doing so would require an ad hoc claim that stronger task sets must be imposed for tasks whose rules are relatively complex, thereby causing their residual disruptive influence on the performance of other subsequent tasks to be greater. In addition, to sustain the TSI hypothesis, it would have to be claimed that the stage at which TSI has a disruptive effect during a subsequent task is temporally and functionally separate from stimulus identification. However, be this as it may, the TSI hypothesis would have considerably more difficulty accounting for the results of our second experiment, which are reported next.

[8] Qualifying this conclusion, inspection of Table 2 reveals that the stimulus-discriminability effect was somewhat smaller on repetitive-task blocks under the high-rule-complexity condition than under the low-rule-complexity condition, although the effect still had the same sign. This decrement presumably occurred because under the high-rule-complexity condition, unlike under the low-rule-complexity condition, nonadjacent values on the relevant and irrelevant perceptual dimensions were used to construct the target and stimulus cards (see Figures 2 and 3). It does not appear that there are necessarily any inherent interactions between the effects of stimulus discriminability and rule complexity on mean RTs or switching-time costs.

Experiment 2

In Experiment 2, to test whether executive control of task switching is mediated jointly by goal shifting and rule activation, we manipulated two factors that should selectively affect the respective durations of these hypothesized stages. The relevant factors were task cuing and rule complexity. According to our model, goal shifting takes less time when explicit cues are available to indicate what the next task is, and rule activation takes more time when performing the next task requires rules that have relatively high complexity (cf. Experiment 1). Furthermore, if the model's assumptions are correct, then the combined effects of these factors on mean switching-time costs should be approximately additive (cf. Sternberg, 1969).

Inspired by some previous studies of task switching (e.g., Jersild, 1927; Spector & Biederman, 1976), Experiment 2 involved tasks that required solving arithmetic problems with alternative numerical combination rules. Rule complexity was manipulated by having participants use addition and subtraction (i.e., low complexity) or multiplication and division (i.e., high complexity) operations. Task cuing was manipulated by either presenting or withholding visual arithmetic-operation signs on both repetitive-task and alternating-task trial blocks. We chose arithmetic for the present task domain because displaying task cues (i.e., the symbols +, −, ×, and ÷) in this context is relatively natural and unobtrusive, so having to encode these cues presumably imposes little extra load on task processes such as stimulus identification. With arithmetic for the task domain, we could also investigate the extent to which our stage model of task switching generalizes beyond classifying visual patterns of geometric objects.

Method

Participants. The participants were 36 undergraduates who belonged to the same population as in Experiment 1 but had not been tested before. Eighteen of them received arithmetic problems with task cues, and the other 18 received arithmetic problems without task cues. Each participant was tested individually during a 1-hr session. None committed more than five errors per trial block.

Laboratory apparatus and environment. The laboratory apparatus and environment were the same as in Experiment 1. During each trial block, the participant solved arithmetic problems on stimulus cards in a deck. This required holding the cards face down, turning them face up one after another, and verbally reporting the numerical solution to a problem on each card. The participant's manual movements were similar to those in Experiment 1.

Stimuli. Task stimuli were constructed with 14-cm × 11.5-cm index cards laminated in clear plastic. There were 12 stimulus cards per deck. Near the center of each card were a two-digit number on the left and a one-digit number on the right (e.g., 56 7). The digits of these numbers were black on a white background and stood about 1.25 cm tall.

Task-cuing conditions. For the cues-present condition, an arithmetic-operation sign (+, −, ×, or ÷) was inserted between the left two-digit and right one-digit numbers on each stimulus card (e.g., 56 + 7). The operation sign indicated the problem type. For the cues-absent condition, there were no operation signs on the cards.

Rule-complexity conditions. There were two rule-complexity conditions: low and high. The low-complexity condition required solving addition and subtraction problems. The high-complexity condition required solving multiplication and division problems. We deemed multiplication to have higher complexity because it involves using the rules of addition together with other supplementary ones. Similarly, we deemed division to have higher complexity because it involves using the rules of subtraction together with other supplementary ones.

Three decks that each contained 12 stimulus cards were created for the low-complexity condition. They were assigned in a counterbalanced fashion to repetitive-task blocks that involved addition, repetitive-task blocks that involved subtraction, and alternating-task blocks that involved switching back and forth between addition and subtraction. Each deck contained a different set of number pairs. When addition was performed, 83% of the problems required carrying a digit from the ones to the tens decimal column; when subtraction was performed, 83% of the problems required borrowing from the tens column. The two-digit numbers ranged from 12 to 68 and were never integral multiples of 10.

Three more decks that each contained 12 stimulus cards were created for the high-complexity condition. They were assigned in a counterbalanced fashion to repetitive-task blocks that involved multiplication, repetitive-task blocks that involved division, and alternating-task blocks that involved switching back and forth between multiplication and division. Each deck contained a different set of number pairs. When multiplication was performed, 83% of the problems required carrying a digit from the ones to the tens decimal column; when division was performed, 83% of the problems required carrying a remainder from the tens to the ones column. Solutions to all of the division problems were integers. The two-digit numbers on the stimulus cards ranged from 36 to 98 and were never an integral multiple of 10.

Design. There were six blocks of test trials per participant: two repetitive-task blocks with the low-complexity tasks (one block for addition and one for subtraction), one alternating-task block with the low-complexity tasks, two repetitive-task blocks with the high-complexity tasks (one block for multiplication and one for division), and one alternating-task block with the high-complexity tasks. The serial orders of the different block types, rule-complexity conditions, and stimulus-card decks were counterbalanced across participants with a nested Latin square design.

Procedure. The procedure was similar to that of Experiment 1. At the start of each session, the participant received instructions and practice. The experimenter informed the participant that he or she would be solving elementary addition, subtraction, multiplication, and division problems. The participant was shown some practice stimulus cards and was told that there would be a two-digit number on the left and a one-digit number on the right of each card. The participant was also told that, for subtraction and division problems, the number on the right should always be subtracted (divided) from (into) the number on the left. Then responses for four different decks of practice stimulus cards with six cards per deck were made.

After practice, the blocks of test trials commenced. At the start of each block, the participant was instructed about what the task or tasks would be. In the cues-absent condition, the experimenter emphasized that the participant should remember which arithmetic operation to perform next. In the cues-present condition, the experimenter emphasized that an arithmetic-operation sign would appear on each stimulus card and that it should be used as a reminder about what the task was. Participants were told to solve the arithmetic problems and verbally report their solutions as quickly as possible without making many errors. Total completion time on each trial block was recorded. Timing began when the experimenter said "Start" and ended when the participant verbally reported the answer to the last arithmetic problem. The experimenter also recorded any errors that occurred along the way.

Results

Table 3 shows mean RTs, error rates, and switching-time costs as a function of trial-block type, rule complexity, and task cuing for each type of arithmetic problem in Experiment 2. Error rates

Table 3
Results of Experiment 2

Trial-block type	Rule complexity	Task cuing	Arithmetic problem type	Mean RT (ms)	Error rate (%)	Switching-time cost (ms)
Repetitive	Low	Present	Addition	2,672	2.3	
	Low	Present	Subtraction	3,528	5.6	
	Low	Absent	Addition	3,125	5.1	
	Low	Absent	Subtraction	3,694	4.2	
	High	Present	Multiplication	6,745	11.1	
	High	Present	Division	5,259	10.7	
	High	Absent	Multiplication	7,866	10.2	
	High	Absent	Division	7,537	10.7	
Alternating	Low	Present	Addition or subtraction	3,158	6.9	63
	Low	Absent	Addition or subtraction	3,875	9.8	508
	High	Present	Multiplication or division	6,713	10.7	776
	High	Absent	Multiplication or division	8,861	13.0	1,265

Note. Mean reaction times (RTs) and error rates were calculated by dividing the number of stimulus cards per deck into the total completion times and errors per trial block. Mean switching-time costs were calculated with Equation 1.

were low on average (8.8% across conditions) and correlated positively with mean RTs ($r = .86$, $p < .001$). The reliability of factor effects on mean RTs and switching-time costs was evaluated as in Experiment 1.

Effects of trial-block type. Trial-block type affected mean RTs reliably, $t(18) = 4.68$, $p < .0005$. On average, participants took longer to respond during alternating-task blocks than during repetitive-task blocks of arithmetic problems, manifesting substantial switching-time costs ($M = 653$ ms, $SE = 140$ ms). As our stage model of task switching predicts, these costs are consistent with component durations being contributed by the executive control processes of goal shifting and rule activation.

Effects of rule complexity. Rule complexity affected mean RTs during both repetitive-task and alternating-task blocks (see Figure 5, left panel). On average, responses were slower for multiplication and division problems than for addition and subtraction problems (mean difference = 3,934 ms, $SE = 425$ ms), $t(18) = 9.26$, $p < .0001$. The magnitude of this effect on mean RTs during repetitive-task blocks (viz., 3,597 ms) presumably manifests how much longer arithmetic calculation and response selection took with high-complexity rules than with low-complexity rules. For now, we assume that an equivalent lengthening of these task processes occurred during alternating-task blocks, which is supported subsequently by further evidence that task and executive control processes may be functionally dissociable from each other.

Nevertheless, during alternating-task blocks, rule complexity had a larger total effect. As a result, switching-time costs were reliably greater for multiplication and division problems than for

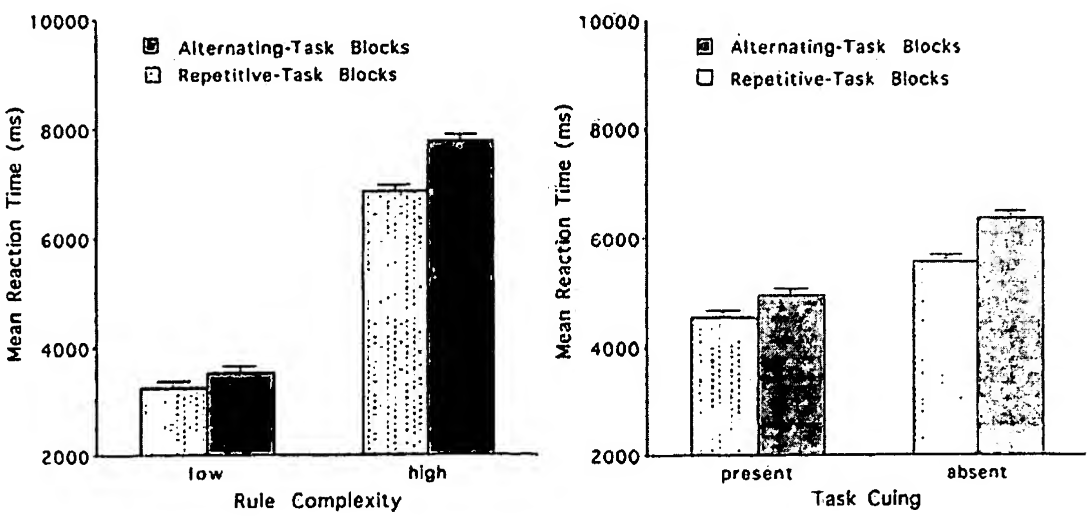

Figure 5. Results of Experiment 2. Left: Mean reaction times (RTs) as a function of rule complexity and trial-block type. Standard errors (lines extending above vertical bars) are based on the interaction among block type, rule complexity, and participants. Right: Mean RTs as a function of task cuing and trial-block type. Standard errors are based on the interaction among block type, task cuing, and participants.

addition and subtraction problems (mean difference = 674 ms, SE = 242 ms), $t(18) = 2.79$, $p < .015$. Again, this difference suggests the existence of a rule-activation stage in executive control that enables task processes such as arithmetic calculation and response selection.

Effects of task cuing. Task cuing also affected mean RTs during both repetitive-task and alternating-task blocks (see Figure 5, right panel). On average, faster responses to arithmetic problems occurred when operation signs were present rather than absent (mean difference = 1,219 ms, SE = 752 ms), $t(18) = 1.62$, p = .061 (one-tailed). The magnitude of this effect during repetitive-task blocks (viz., 1,005 ms) presumably manifests how much less time arithmetic calculation and response-selection took with task cues, and as mentioned already, we assume that an equivalent shortening of these task processes occurred during alternating-task blocks.

Consistent with the latter assumption, the task-cuing and rule-complexity effects on mean RTs interacted during both repetitive-task and alternating-task blocks. In particular, task cues shortened RTs more when high-complexity (multiplication and division) rules were involved (Figure 6, left panel; mean difference = 1,411 ms, SE = 850 ms), $t(18) = 1.66$, p = .06 (one-tailed). This interaction is what we would expect if, as suggested before, both the cuing and complexity factors influenced a task process such as arithmetic calculation or response selection. Furthermore, the magnitude of this interaction was approximately the same regardless of trial-block type (1,390 ms and 1,431 ms for repetitive-task and alternating-task blocks, respectively), suggesting that it did not involve executive control processes per se.

Nevertheless, during alternating-task blocks, task cuing had a larger total effect on mean RTs than it did during repetitive-task blocks. As a result, switching-time costs were reliably lower with task cues present rather than absent (see Figure 6, right panel; mean difference = 467 ms, SE = 278 ms), $t(18) = 1.68$, p = .056 (one-tailed). This establishes the existence of a goal-shifting stage in executive control that complements rule activation. As predicted by our model, goal shifting apparently makes its own distinct contribution to task switching, and its duration depends on the extent to which explicit information is available about the next type of task that has to be performed.

Additivity of task-cuing and rule-complexity effects on switching-time costs. Also, as our model predicted, the effects of task cuing and rule complexity on mean switching-time costs were approximately additive (see Figure 6, right panel; mean interaction = 45 ms, SE = 536 ms), $t(18) = 0.084$, $p > .5$. This additivity supports our claim that goal shifting and rule activation are temporally separate and selectively influenced stages of executive control in task switching. It is especially interesting that such selectivity occurred here even though task cuing and rule complexity both influenced at least one task process (e.g., arithmetic calculation) in common, as manifested by the interaction between the effects of these factors on mean RTs (see Figure 6, left panel). The present additive factor effects on mean switching-time costs combined with the accompanying interactive factor effects on mean RTs supplement the previous pattern of results regarding rule-complexity and stimulus-discriminability effects in Experiment 1, which likewise showed that executive control processes may be separable from the task processes enabled by them.

Discussion

The results of Experiment 2 substantially augment those of Experiment 1, reinforcing our stage model of executive control in task switching. In a new task domain (i.e., arithmetic problem

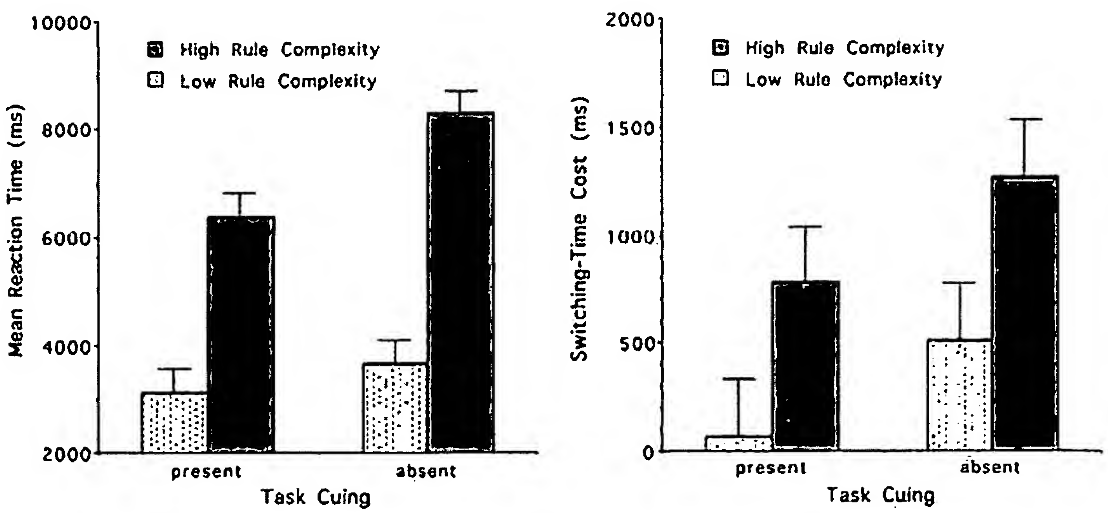

Figure 6. Additional results of Experiment 2. Left: Mean reaction times as a function of task cuing and rule complexity. Standard errors (lines extending above vertical bars) are based on the interaction among task cuing, rule complexity, and participants. Right: Mean switching-time costs as a function of task cuing and rule complexity. Standard errors are based on the interaction among trial-block type, task cuing, rule complexity, and participants.

solving), participants had longer mean RTs during alternating-task blocks than during repetitive-task blocks. Switching-time costs were again reliably greater for high-complexity than low-complexity tasks, suggesting a rule-activation stage of executive control. Furthermore, greater switching-time costs occurred when task cues were absent rather than present, which may be attributed to a complementary goal-shifting stage of executive control whose duration depends on whether or not explicit information is available about the next type of task to be performed. The task-cuing and rule-complexity effects on mean switching-time costs were approximately additive, as should occur if goal shifting and rule activation are temporally discrete stages of processing. That these executive control processes may be functionally separate from concomitant task processes also appeared evident from other aspects of the present data.

In contrast, the TSI hypothesis has considerable trouble accounting for these results. As formulated to date, this hypothesis does not postulate any distinct well-specified component processes that jointly contribute to the time costs of task switching but that can be influenced selectively by different factors. Thus, it is not clear how additive-factor effects on mean switching-time costs such as those manifested in Experiment 2 could stem from TSI per se. Instead, such effects seem more consistent with our stage model of executive control in task switching.

However, one result of Experiment 2 poses a potential puzzle for our model. With respect to switching between addition and subtraction problems accompanied by task cues, we found a mean time cost of only 63 ms (see Figure 6, right panel). This raises a question about the extent to which executive control was needed when these problems were involved. If goal shifting and rule activation constitute crucial operations, then how could their combined durations have yielded such a small switching-time cost?

One possible answer is that on alternating-task blocks of Experiment 2, executive control processes for switching to the next arithmetic task overlapped temporally with physical movements for completing the prior task. Such overlap may have occurred while participants placed the preceding task's stimulus card on the table and began turning over the next task's card. This could spuriously reduce estimated switching-time costs, attenuating the apparent contributions of goal shifting and rule activation so that they seem especially small for relatively easy addition and subtraction problems (cf. Rogers & Monsell, 1995).

A second possibility is that addition problems accompanied by plus signs have special status. Perhaps the production rules used in solving these problems are permanently enabled in procedural long-term memory so that rule activation need not contribute to switching-time costs for them. This seems plausible because arithmetic addition—like reading printed words—is learned at an early age and practiced regularly throughout many people's lives.

However, these preceding possibilities are speculative. Thus, we tried to check them further during a third experiment. Experiment 3 also provided a test of some further predictions on the basis of our stage model of executive control.

Experiment 3

In Experiment 3, we focused on one additional important prediction about the executive control of task switching. According to our stage model, the time taken for switching from a first task A to a second task B will not necessarily equal the time taken for switching in the opposite direction. Instead, there may be systematic differences between switching-time costs, depending on the direction of task switching.

This prediction follows from the model because, under it, task switching entails a rule-activation stage of executive control whose duration may differ as a function of the preceding and following tasks' rules. More precisely, we hypothesize that greater switching-time costs may tend to occur when people switch from a relatively familiar task to a relatively unfamiliar task rather than vice versa (cf. Monsell et al., 2000). The present *task-familiarity hypothesis* could hold because the rule-activation stage takes more time to enable the rule set for a current unfamiliar task or because it takes more time to disable the rule set for an immediately prior familiar task (cf. Goschke, 2000; Mayr & Keele, 2000; Mayr & Kliegl, 2000). Thus, demonstrating such asymmetric familiarity-dependent time costs of task switching would provide more evidence for a specific rule-activation stage and for our model's basic assumptions. Doing so would also cast more doubt on the TSI hypothesis. This follows because, on the basis of hypothesized TSI, Allport et al. (1994) have claimed that switching from relatively unfamiliar tasks (e.g., ink-color naming) to relatively familiar tasks (e.g., color-word reading) takes more time than does switching from familiar to unfamiliar tasks, which is exactly opposite to the asymmetry expected from our task-familiarity hypothesis.

To test for asymmetric task switching and to assess the task-familiarity hypothesis about the asymmetry's source, Experiment 3 entailed again having participants solve various types of arithmetic problems during alternating-task and repetitive-task blocks. As in Experiment 2, there were two rule-complexity conditions: low (addition or subtraction) and high (multiplication or division). Given plausible supplementary assumptions, each of these conditions yielded mean RTs and switching-time costs relevant to our objectives.

Specifically, we assume that among the tasks with low rule complexity, addition is relatively more familiar than subtraction. This assumption is justified on several grounds. For example, in elementary mathematics education, the topic of addition has traditionally been introduced and practiced before subtraction. Also, solving multicolumn subtraction problems—which include steps of between-columns "borrowing"—often entails addition as well, whereas solving multicolumn addition problems does not entail subtraction. Consequently, addition would receive more practice and thereby become more familiar than subtraction. Because of similar considerations, we assume that among the tasks with high rule complexity, multiplication is relatively more familiar than division. This assumption is justified given that, in elementary mathematics education, the topic of multiplication has traditionally been introduced and practiced before division. Also, solving multicolumn division problems often entails multiplication, whereas solving multicolumn multiplication problems does not entail division. Consequently, multiplication would receive more practice and thereby become more familiar than division. In turn, this latter difference would contribute further to addition being more familiar than subtraction, because solving multicolumn multiplication

problems requires addition, whereas solving multicolumn division problems requires subtraction.

Supplemented by these assumptions, the task-familiarity hypothesis and our stage model of executive control led us to expect two related results from Experiment 3. First, the mean time cost of switching from addition to subtraction problems should exceed the mean time cost of switching in the opposite direction. Second, switching from multiplication to division problems should take longer on average than switching in the opposite direction. Insofar as task familiarity and rule complexity both affect the rule-activation stage, one might also expect their effects on mean switching-time costs to interact. The expected directions of the task-familiarity effect and its interaction with rule complexity are not what the TSI hypothesis would predict (cf. Allport et al., 1994). Thus, Experiment 3 allowed us to replicate and extend Experiment 2 regarding rule activation and the mechanisms through which it mediates rule-complexity effects.

Because this extension required separately measuring time costs of task switching in opposite directions, the method of Experiment 3 differed somewhat from that of Experiment 2. Instead of displaying arithmetic problems on index cards, we displayed each problem on a video screen during blocks of discrete trials. Participants responded to each problem by pressing a key to indicate where the correct numerical solution was located in an array of alternative solutions beneath the problem. RT and accuracy of individual responses were measured relative to stimulus onsets, enabling the time cost of task switching in each direction to be precisely measured. Moreover, so that we could focus on the role played by rule activation under these conditions, all arithmetic problems were accompanied by explicit task cues, which facilitated goal shifting and reduced its contribution to switching-time costs, making the contributions of rule activation relatively more salient.

78 X 3

234 375 116 125	72 135 324 504	304 607 288 162	325 637 156 224

Figure 7. Visual display for a representative trial in the high-rule-complexity condition of Experiment 3. At the top is an arithmetic problem (78 × 3) to be solved, and at the bottom are four target boxes with potential numerical solutions to the problem. Here the correct solution is 234, which appears in the left-most target box.

Method

Participants. The participants were 36 undergraduates who belonged to the same population as in the previous experiments but had not been tested before. Each participant was tested individually during a 1-hr session. None committed more than four errors per trial block.

Laboratory apparatus and environment. Each participant sat at a table on which was placed a Macintosh SE/30 personal computer with a 9-in. (23-cm) black-and-white video screen and Apple keyboard. The computer controlled the presentation of stimuli, recording of responses, and storage of data. The video screen was used for displaying task stimuli. The viewing distance to the video screen was approximately 50 cm. The keyboard was placed in front of the video screen. The participant put the index and middle fingers of each hand on the keyboard and pressed the *V, B, N,* or *M* key to make individual responses. Four small labels with the digits *1, 2, 3,* and *4* on them were attached to these keys in left-to-right order, which paralleled the left-to-right order of four target boxes that appeared on the video screen during each trial.

Stimuli. On each trial, an arithmetic problem was displayed above a row of four target boxes (e.g., see Figure 7). The arithmetic problems and target boxes were centered horizontally on the video screen. A two-digit number appeared on the left and a one-digit number appeared on the right of each problem, with an arithmetic-operation sign between them. The digits and symbols for the problem were printed in 36-point bold Geneva font and were about 1 cm tall. The vertical distance from the top of the video screen to the center of each arithmetic problem was about 5 cm; the horizontal distance between the centers of the left and right numbers in the problem was about 4 cm. The vertical distance between the centers of each arithmetic problem and row of target boxes was about 5 cm. The four target boxes had 4-cm sides. The horizontal distance between the outer edges of the far left and far right target boxes was about 17 cm. Near its corners, each target box contained four different (typically multidigit) numbers that were printed in 24-point bold Geneva font and were about 0.6 cm tall. One of them was the correct solution to the problem displayed above.

Design. The design of the sessions and trial blocks was similar to that of Experiment 2. Each session had an instruction phase, four practice trial blocks, and six test trial blocks. There were 6 trials per practice block. Included in each test block were 2 warm-up trials followed by 12 test trials. Three test blocks involved the tasks with low-complexity rules: one repetitive-task block for addition, one repetitive-task block for subtraction, and one alternating-task block for addition and subtraction. The other three test blocks involved the tasks with high-complexity rules: one repetitive-task block for multiplication, one repetitive-task block for division, and one alternating-task block for multiplication and division.

During the test trial blocks, the arithmetic problems were the same as in Experiment 2. Their order of presentation was randomized. The contents of the target boxes remained constant throughout a block but changed between blocks. The relative frequencies with which the correct solutions to the problems appeared in the various target boxes were equated. Across participants, the serial orders of the rule-complexity conditions and trial-block types were also counterbalanced with a nested Latin square design.

Procedure. At the start of each session, the experimenter instructed the participant about the various tasks and the arithmetic problems for them. The participant watched the experimenter demonstrate the procedure with the video screen and keyboard. Then the practice and test trial blocks commenced.

Before each trial block, information was displayed on the video screen about what the forthcoming task or tasks would be, and the participant was encouraged to use the task cues for keeping track of which task had to be performed next. After the participant finished reading this information, the successive trials of the block followed, with a visual display being presented on each trial as in Figure 7. For each display, the participant looked at the arithmetic problem therein, solved it with the relevant arithmetic operation, and then found which target box contained this solution. The participant indicated the location of this box by pressing the response key that corresponded to it. RT was measured from stimulus onset until a keypress occurred, and response accuracy was also recorded. Prior instructions encouraged the participant to respond as quickly as possible on each

Table 4
Results of Experiment 3

Trial-block type	Rule complexity	Arithmetic-problem type: Prior task	Arithmetic-problem type: Current task	Mean RT (ms)	Error rate (%)	Switching-time cost (ms)
Repetitive	Low	Addition	Addition	4,371	5.1	
	Low	Subtraction	Subtraction	4,805	5.3	
	High	Multiplication	Multiplication	8,541	8.8	
	High	Division	Division	6,813	10.2	
Alternating	Low	Subtraction	Addition	4,400	5.6	29
	Low	Addition	Subtraction	5,044	4.2	239
	High	Division	Multiplication	9,479	6.5	938
	High	Multiplication	Division	8,063	7.9	1,250

Note. Mean switching-time costs were calculated by subtracting mean reaction times (RTs) on repetitive-task blocks from mean RTs on alternating-task blocks. *Prior task* refers to the task performed on the immediately preceding trial, and *current task* refers to the task for which data are reported.

trial without making many errors. The RSI between trials was approximately 150 ms.

At the end of each trial block, summary feedback was displayed to emphasize both speed and accuracy of performance. The feedback indicated mean RT and total number of errors for the block. If no errors occurred during the block, the participant was also informed "You are doing very well."

Results

Table 4 shows mean RTs, error rates, and switching-time costs as a function of trial-block type, rule complexity, and direction of task switching for each type of arithmetic problem in Experiment 3.[9] Error rates were low on average (6.7% across conditions) and correlated positively with mean RTs ($r = .57$, $p < .1$, one-tailed). The reliability of factor effects on mean RTs and switching-time costs was evaluated as in the previous experiments.

Effects of trial-block type. Trial-block type affected mean RTs reliably, $t(18) = 2.59$, $p < .05$. On average, participants took longer to respond during alternating-task blocks than during repetitive-task blocks, manifesting substantial switching-time costs ($M = 614$ ms, $SE = 237$ ms). This replicated the results of Experiment 2, in which the overall mean switching-time cost was 653 ms for arithmetic problems such as the present ones. In most respects, the present time costs were consistent with component durations being contributed by the executive control processes of goal shifting and rule activation, although the time taken for switching from subtraction to addition problems was again very short ($M = 29$ ms).

Effects of rule complexity. Rule complexity affected mean RTs during both repetitive-task and alternating-task blocks (see Figure 8, left panel). As before (Experiment 2), responses were slower on average for multiplication and division problems than for addition and subtraction problems (mean difference = 3,569 ms, $SE = 298$ ms), $t(18) = 12.0$, $p < .0001$. The magnitude of this effect on mean RTs during repetitive-task blocks (viz., 3,089 ms) presumably manifests how much longer arithmetic calculation and response selection took with high-complexity rules than with low-complexity rules. For reasons already mentioned, we again assume that an equivalent lengthening of these task processes occurred during alternating-task blocks.

Nevertheless, during alternating-task blocks, rule complexity again had a larger total effect. As a result, switching-time costs were reliably greater for multiplication and division problems than for addition and subtraction problems (mean difference = 960 ms, $SE = 355$ ms), $t(18) = 2.70$, $p < .05$. This helps confirm the existence of a rule-activation stage in executive control that enables task processes such as arithmetic calculation and response selection.

Asymmetry of task switching. For both the arithmetic problems that involved low-complexity rules and the arithmetic problems that involved high-complexity rules, there tended to be an asymmetry in task switching (see Figure 8, right panel), consistent with our hypothesis about the nature of rule activation in executive control. On average, switching from addition to subtraction took longer than switching from subtraction to addition (mean difference = 210 ms). Furthermore, on average, switching from multiplication to division took longer than switching from division to multiplication (mean difference = 312 ms). These differences are what we would expect in terms of the present task-familiarity hypothesis, given that addition and multiplication problems are presumably more familiar than subtraction and division problems, respectively.

However, the results of Experiment 3 failed to be fully conclusive in this regard. Overall, the relative familiarity and unfamiliarity of the tasks between which participants switched here did not affect switching-time costs reliably; for familiar-to-unfamiliar versus unfamiliar-to-familiar task switching, the mean difference was 261 ms ($SE = 298$ ms), $t(18) = 0.66$, $.20 < p < .30$ (one-tailed). This unreliability stemmed from substantial between-subjects variance in the magnitudes of the familiarity effects. Some possible

[9] A formula different from Equation 1 was used to calculate the mean switching-time costs in Experiment 3. This followed because the procedure of Experiment 3 involved discrete trials, a separate RT was measured on each trial, and the test trials of each block were preceded by similarly arranged warm-up trials. Given this arrangement, we calculated the mean switching-time costs simply by subtracting the mean RTs on repetitive-task blocks for each task from the corresponding mean RTs on alternating-task blocks. For similar reasons, the latter formula was also used to calculate the mean switching-time costs in Experiment 4.

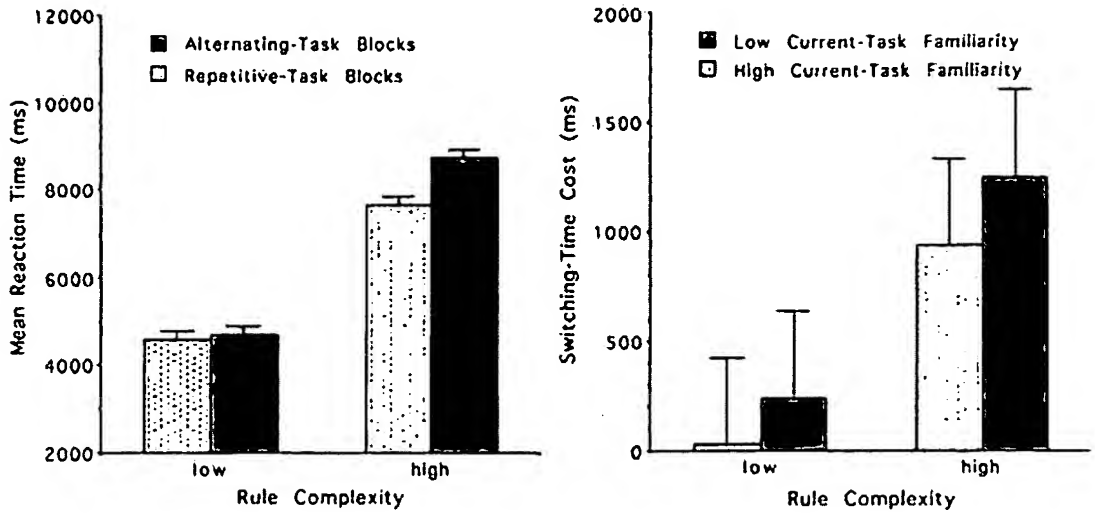

Figure 8. Results of Experiment 3. Left: Mean reaction times (RTs) as a function of rule complexity and trial-block type. Standard errors (lines extending above vertical RT bars) are based on the interaction among block type, rule complexity, and participants. Right: Mean switching-time costs as a function of rule complexity and current-task familiarity. Standard errors are based on the interaction among trial-block type, current-task familiarity, and participants.

causes of these individual differences are considered in the following discussion.

Discussion

Using arithmetic problem-solving tasks once more, Experiment 3 replicated and extended Experiment 2, yielding some results that further confirm predictions based on our stage model of executive control in task switching. Although the present procedure involved discrete RT trials rather than continuous card sorting, mean switching-time costs were again positive and increased reliably with rule complexity, as if they stemmed from a rule-activation stage of executive control. We also found that these costs tended to depend on the direction of task switching, being larger on average when participants switched from putatively more familiar (i.e., addition or multiplication) to less familiar (i.e., subtraction or division) tasks rather than vice versa. This overall pattern held both for switching between tasks that involved low rule complexity and for switching between tasks that involved high rule complexity. Thus, consistent with our task-familiarity hypothesis, it appears in some respects that perhaps the rule-activation stage takes longer to enable the rules of unfamiliar tasks or to disable the rules of familiar tasks.

Concomitantly, we obtained more evidence against the TSI hypothesis. As proposed by Allport et al. (1994), the TSI hypothesis implies that, on average, switching from more familiar (e.g., color-word reading) to less familiar (e.g., ink-color naming) tasks should take less time than switching in the opposite direction. Nevertheless, during Experiment 3, familiar-to-unfamiliar task switching tended to take more time than unfamiliar-to-familiar switching. It may therefore be concluded that TSI is not a general salient phenomenon in task switching or that the dynamics of TSI do not consistently conform to Allport et al.'s (1994) claims. Instead, under at least some conditions, contributions by separable executive control processes to switching-time costs appear to predominate.

Yet, in some other respects, Experiment 3 yielded inconclusive data. Although consistent with the task-familiarity hypotheses, the overall asymmetry of task switching found here did not reach a compelling level of reliability. There was considerable between-subjects variance in the prior-task and current-task familiarity effects on mean switching-time costs. One possible source of this variability could be systematic individual differences in participants' past mathematics and science training. Because of such differences, our initial assumptions that addition problems are significantly more familiar than subtraction problems and that multiplication problems are significantly more familiar than division problems may not hold across all participants. If so, then perhaps the task-familiarity hypothesis is still tenable, but further research is needed to test it under more relevant conditions.

Another problematic result of Experiment 3, similar to what occurred in Experiment 2, concerns the time cost of switching from subtraction to addition problems accompanied by plus signs. We found that, on average, this switching-time cost was only 29 ms, even though there was a very short RSI between problems. The small magnitude of this mean time cost suggests that switching to an explicitly cued and extremely familiar task such as arithmetic addition may not always require rule activation per se. Instead, perhaps the rules for solving signed addition problems—like the rules for reading familiar printed words—are permanently enabled in procedural long-term memory, thereby requiring the rule-activation stage of executive control to take little or no extra time for fully enabling them. In light of these considerations, we

sought to test our stage model and the task-familiarity hypothesis during a fourth experiment whose design was perhaps more suitable for evaluating their potential merits.

Experiment 4

In Experiment 4, to characterize the rule-activation stage of executive control more fully and to evaluate the task-familiarity hypothesis further, we returned to having participants alternate between pattern-classification tasks. As before (Experiment 1), these tasks required either low-complexity (unidimensional) or high-complexity (bidimensional) rules to classify visual patterns with respect to their shapes, sizes, colors, and numerosities of geometric objects. For each task, we measured individual RTs on discrete trials of alternating-task and repetitive-task blocks, and we calculated the time costs of switching between tasks in one direction versus another, following the same steps taken in Experiment 3. This helped set the stage for again examining whether differences in task familiarity yield systematic asymmetries in task switching.

Our approach to analyzing and interpreting the results of Experiment 4 was also inspired by findings that we obtained during a related subsidiary study. In that study, another group of participants made a series of paired-comparison familiarity judgments about the perceptual dimensions of the present pattern-classification tasks. For each judgment, a participant received a pair of perceptual-dimension names chosen from the ones that Experiment 4 involved. The participant indicated which member of the pair corresponded to the perceptual dimension about which he or she more frequently made conscious visual decisions during daily routines. These familiarity judgments were made for all possible pairs of the four relevant perceptual dimensions. From them, we constructed each participant's rank order of the dimensions' subjective familiarity (4 = *most familiar*, 1 = *least familiar*) based on standard psychological-scaling methods (Coombs, 1964). The ranks for each perceptual dimension were then averaged across participants, yielding a numerical scale of subjective-familiarity scores. On this scale, the scores of the shape, size, numerosity, and shading dimensions were, respectively, 3.14, 2.71, 2.71, and 1.43, with larger numbers representing higher subjective familiarity (maximum score = 4, minimum score = 1).[10]

Supplemented by these findings, the task-familiarity hypothesis implies that certain asymmetries in task switching should emerge during Experiment 4. Specifically, switching from the shape-classification to the numerosity-classification task should take longer than switching in the opposite direction, because making decisions about object numerosities is putatively less familiar than making decisions about object shapes (viz., 2.71 vs. 3.14 on our subjective-familiarity scale). Similarly, switching from the size-classification to the shading-classification task should take longer than switching in the opposite direction, because making decisions about object shading is putatively less familiar than making decisions about object sizes (viz., 1.43 vs. 2.71 on our subjective-familiarity scale). By demonstrating the occurrence of these expected differences in mean switching-time costs, Experiment 4 supports the task-familiarity hypothesis while casting more doubt on the general veracity of the TSI hypothesis, which makes predictions opposite to what we expect here. Also, our results let us learn more about rule activation in executive control, because the design of Experiment 4 provided a way to separate the effects of prior-task familiarity from the effects of current-task familiarity on this stage. Through this separation, we can see whether mean switching-time costs are greater because it takes longer to disable the rules of prior familiar tasks or because it takes longer to enable the rules of current unfamiliar tasks.

Method

Participants. The participants were 24 undergraduates who belonged to the same population as in the previous experiments but had not been tested before. Each participant was tested individually during a 1-hr session.

Laboratory apparatus and environment. A personal computer controlled the presentation of stimuli, recording of responses, and storage of data. The computer's video screen was used for displaying task stimuli, which were created with the Macintosh Superpaint utility program. The participant placed the index and middle fingers of each hand on the computer's keyboard and pressed the *V*, *B*, *N*, or *M* key to make individual responses. Four small labels with the digits *1*, *2*, *3*, and *4* on them were attached to these keys in left-to-right order, which paralleled the left-to-right order of four target boxes that appeared on the video screen during each trial. Other aspects of the laboratory apparatus and environment were as in Experiment 3.

Stimuli. For each trial, the computer presented a visual display that contained a task cue, probe stimulus, and row of four targets (see Figure 9). The cue, stimulus, and targets were centered horizontally on the video screen. The stimulus and targets appeared in square boxes whose edges were 4 cm long. The horizontal distance between the outer edges of the far left and far right target boxes was 17 cm. The vertical distance between the centers of the task cue and task stimulus was 2.7 cm. The vertical distance between the centers of the task stimulus and row of targets was 4.5 cm. The cues were printed in uppercase 14-point bold Geneva font.

Within the boxes for the task stimuli and targets were patterns of geometric objects similar to those of Experiment 1. Each pattern was defined by a combination of four perceptual dimensions (shape, numerosity, size, and background shading). Across stimuli, the objects had four different shapes (triangle, circle, star, and cross), four levels of numerosity (1, 2, 3, and 4), four sizes (approximately 3, 8, 13, and 18 mm in height), and four levels of background shading (dark, medium, light, and white). The background rather than interior shading of the objects was manipulated because this helped reduce possible confoundings with the objects' sizes.

Low-rule-complexity condition. For the low-rule-complexity condition, there were four different tasks that required classifying stimuli with respect to either object shape, size, shading, or numerosity. The cues for these tasks were, respectively, the words SHAPE, SIZE, SHADING, and

[10] Seven students at the University of Michigan contributed to these scores. Every participant was transitive in his or her paired-comparison judgments (i.e., if a participant judged dimension *A* to be more familiar than dimension *B*, and dimension *B* to be more familiar than dimension *C*, then he or she always judged dimension *A* to be more familiar than dimension *C*). The coefficient of concordance between the participants' rank orders of the dimensions' familiarities was reliably positive, $W = 0.33$, $S(5) = 81.0$, $p < .05$, one-tailed (cf. Hays, 1963; Siegel, 1956). For 5 of the 7 participants, the subjective familiarity of the shape dimension ranked either first or second highest; for 6 participants, the subjective familiarity of the shading dimension ranked lowest. However, for 1 atypical participant, the shading dimension had the highest subjective familiarity. Excluding him, the other 6 participants had a markedly greater coefficient of concordance, $W = 0.61$, $S(4) = 109.8$, $p < .01$.

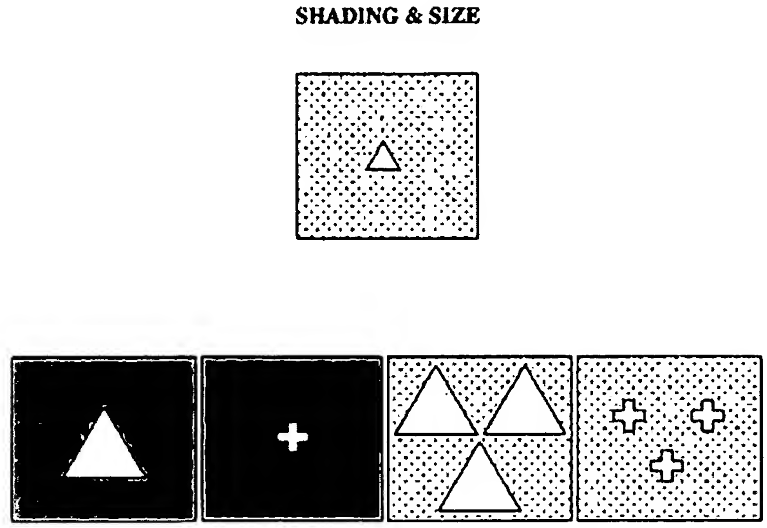

Figure 9. Visual display for a representative trial in the high-rule-complexity condition of Experiment 3. The task cue (SHADING & SIZE), task stimulus (one small white triangle with light background), and targets appear respectively at the top, middle, and bottom of the display.

NUMBER. The same set of four targets was used in each case. They were constructed such that no two of them matched each other on any dimension. Depending on which task was involved, a particular stimulus could correctly match any one of the four targets. Participants performed each task either singly on repetitive-task trial blocks or in combination with another low-complexity task on alternating-task trial blocks. Within each block, equal numbers of stimuli matched the relevant perceptual dimension with respect to each of its four values. Also, each stimulus had values on three irrelevant dimensions that matched those of the other targets.

High-rule-complexity condition. For the high-rule-complexity condition, there were two different tasks that required classifying stimuli with respect to either object shape and number or object size and background shading. The cues for these tasks were, respectively, the phrases SHAPE AND NUMBER and SIZE AND SHADING. Participants performed each task either singly on repetitive-task blocks or in combination with the other task on alternating-task blocks. In each case, there again were four targets arranged as in Experiment 1 (cf. Figures 3 and 9).

Design. The design of the sessions and trial blocks was similar to that of Experiment 3. Each session had an instruction phase, 8 practice trial blocks, and 12 test trial blocks. There were 12 trials per practice block. Included in each test block were 4 warm-up trials followed by 24 test trials. The test trials were arranged to form 4 repetitive-task blocks with the low-complexity rules (1 block for each perceptual dimension), 2 alternating-task blocks with the low-complexity rules (1 block for switching between shape and numerosity and another for switching between size and shading), 2 repetitive-task blocks with the high-complexity rules that involved conjunctions of shape and numerosity, 2 repetitive-task blocks with the high-complexity rules that involved conjunctions of size and shading, and 2 alternating-task blocks with the high-complexity rules. Across participants, the serial orders of the relevant perceptual dimensions, trial-block types, and rule complexity were counterbalanced with a nested Latin square design. Within each trial block, the order of stimulus presentation was randomized.

Procedure. The procedure was similar to that of Experiment 3. At the start of each session, the experimenter explained the four perceptual dimensions to the participant. The four low-complexity and two high-complexity classification rules were introduced. The participant watched the experimenter demonstrate the procedure with the video screen and keyboard. Then the practice and test trial blocks commenced.

Before each trial block, information was displayed on the video screen about what the forthcoming task or tasks would be, and the participant was encouraged to use the task cues for keeping track of which task had to be performed next. After the participant finished reading this information, the successive trials of the block followed, with a visual display being presented on each trial as in Figure 9. For each display, the participant looked at the task stimulus and decided which of the four target stimuli below it was matched according to the relevant classification rule. Then the participant pressed the response key corresponding to the box that contained this target. RT was measured from stimulus onset until a keypress occurred, and response accuracy was also recorded. Prior instructions encouraged the participant to respond as quickly as possible on each trial without making many errors. The RSI between trials was approximately 150 ms.

At the end of each trial block, summary feedback was displayed to emphasize both speed and accuracy of performance. The feedback indicated the mean RT and total number of errors for the block. If no errors occurred during the block, the participant was also informed "You are doing very well."

Results

Table 5 shows mean RTs, error rates, and switching-time costs as a function of trial-block type, rule complexity, and direction of task switching for each visual pattern-classification task in Experiment 4. Error rates were low on average (5.7% across conditions) and correlated positively with mean RTs ($r = .94$, $p < .001$). The reliability of factor effects on mean RTs and switching-time costs was evaluated as in the previous experiments.

Effects of trial-block type. Trial-block type affected mean RTs reliably, $t(12) = 26.2$, $p < .0001$. On average, participants took longer to respond during alternating-task blocks than during repetitive-task blocks, manifesting substantial switching-time costs ($M = 1{,}096$ ms, $SE = 42$ ms). This replicated the results of

Table 5
Results of Experiment 4

Trial-block type	Rule complexity	Relevant perceptual dimensions: Prior task	Relevant perceptual dimensions: Current task	Mean RT (ms)	Error rate (%)	Switching-time cost (ms)
Repetitive	Low	Size	Size	1,030	3.8	
	Low	Shading	Shading	705	3.0	
	Low	Shape	Shape	865	2.8	
	Low	Numerosity	Numerosity	693	2.3	
	High	Size and shading	Size and shading	1,038	3.3	
	High	Shape and numerosity	Shape and numerosity	1,108	4.0	
Alternating	Low	Shading	Size	1,532	6.6	502
	Low	Size	Shading	1,380	3.1	675
	Low	Numerosity	Shape	1,357	4.5	492
	Low	Shape	Numerosity	1,462	3.1	769
	High	Shape and numerosity	Size and shading	2,468	9.9	1,430
	High	Size and shading	Shape and numerosity	2,841	13.7	1,733

Note. Mean switching-time costs were calculated by subtracting mean reaction times (RTs) on repetitive-task blocks from mean RTs on alternating-task blocks. *Prior task* refers to the task performed on the immediately preceding trial, and *current task* refers to the task for which data are reported.

Experiment 1, in which the overall mean switching-time cost was 975 ms for visual pattern-classification tasks. Again, it appears that as our stage model of executive control assumes, task switching was probably mediated by goal-shifting and rule-activation stages.

Effects of rule complexity. Rule complexity affected mean RTs during both repetitive-task and alternating-task blocks (see Figure 10, left panel). As before (Experiment 1), high-complexity rules yielded slower pattern-classification responses than low-complexity rules (mean difference = 736 ms, SE = 47 ms), $t(12) = 15.8$, $p < .0001$. The magnitude of this effect on mean RTs during repetitive-task blocks (viz., 250 ms) presumably manifests how much longer judgment and response selection took with high-complexity (i.e., bidimensional) than with low-complexity (i.e., unidimensional) classification rules. For reasons already mentioned, we again assume that an equivalent lengthening of these task processes occurred during alternating-task blocks.

Nevertheless, during alternating-task blocks, rule complexity had a larger total effect. As a result, switching-time costs were reliably greater for the pattern-classification tasks with high-complexity rules than for those with low-complexity rules (mean difference = 972 ms, SE = 51 ms), $t(18) = 19.1$, $p < .0001$. Again, this helps demonstrate the existence of a rule-activation stage in executive control that enables task processes such as response selection.

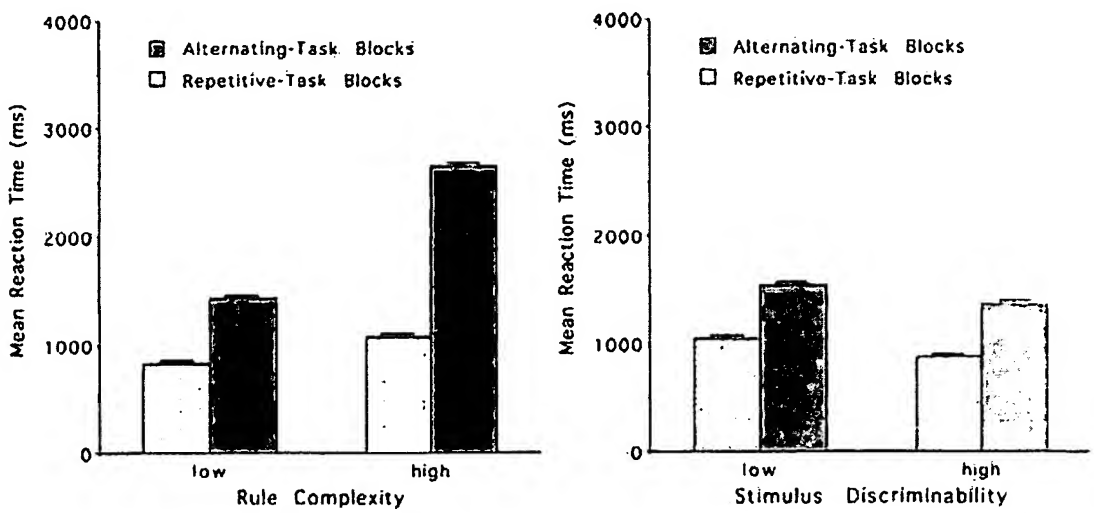

Figure 10. Results of Experiment 4. Left: Mean reaction times (RTs) as a function of rule complexity and trial-block type. Standard errors (lines extending above vertical RT bars) are based on the interaction among block type, rule complexity, and participants. Right: Mean RTs as a function of stimulus discriminability and trial-block type for the size and shape classification tasks. Standard errors are based on the interaction among block type, stimulus discriminability, and participants.

Effects of stimulus discriminability. Also consistent with our stage model and the results of Experiment 1, stimulus discriminability again had approximately additive effects on mean RTs. For example, on repetitive-task blocks, the size-classification task yielded reliably slower responses than the shape-classification task (mean difference = 165 ms, SE = 53 ms), $t(12) = 3.10$, $p < .01$. Yet switching-time costs for these two tasks were about the same on alternating-task blocks (mean difference = 10 ms, SE = 71 ms), $t(12) = 0.14$, $p > .5$. This further confirms our model's prediction that task processes (e.g., stimulus identification and response selection) and executive control processes (e.g., goal shifting and rule activation) can be empirically dissociated and influenced selectively by different experimental factors.[11]

Asymmetry of task switching. As expected from the task-familiarity hypothesis, some systematic asymmetries of task switching occurred during alternating-task blocks with the visual pattern-classification tasks that had low-complexity (unidimensional) rules. Recall that on our scale for subjective familiarity of visual decisions about the perceptual dimensions of these tasks, the shape dimension had the highest score, the numerosity and size dimensions had intermediate scores, and the shading dimension had the lowest score (mean familiarities = 3.14, 2.71, 2.71, and 1.43, respectively). Correspondingly, we found that switching from the shape-classification task to the numerosity-classification task took reliably longer on average than switching in the opposite direction (mean difference = 277 ms, SE = 54 ms), $t(12) = 5.09$, $p < .0001$. Furthermore, switching from the size-classification task to the shading-classification task took reliably longer on average than switching in the opposite direction (mean difference = 173 ms, SE = 45 ms), $t(12) = 3.82$, $p < .005$. These differences support the task-familiarity hypothesis, which predicts that familiar-to-unfamiliar switching should be slower than unfamiliar-to-familiar switching, because the rule-activation stage takes longer to enable the rules of current unfamiliar tasks or to disable the rules of prior familiar tasks.

To precisely quantify the relative contributions of prior-task and current-task familiarity to mean switching-time costs, and thereby to learn more about the nature of the rule-activation stage, we performed a multiple linear regression analysis with two predictor variables and one predicted variable for the visual pattern-classification tasks that had low-complexity rules. The first predictor variable was the subjective-familiarity score of the prior task from which the participants switched; the second predictor variable was the subjective-familiarity score of the current task to which the participants switched. The predicted variable was the mean switching-time cost as a function of the prior-task and current-task familiarities. With these variables, the regression analysis yielded

$$T_S = 487 + 118 \times F_P - 69 \times F_C, \qquad (2)$$

where T_S is mean switching-time cost in milliseconds, F_P is prior-task familiarity, F_C is current-task familiarity, the additive intercept coefficient is in milliseconds, and the multiplicative predictor coefficients are in millisecond/familiarity units.

Equation 1 yielded a large multiple-correlation coefficient (R = .78) and accounted for a reliably positive proportion of the variance in mean switching-time costs, $R^2 = .61$, $t(12) = 2.59$, $p < .05$. The prior-task and current-task familiarity scores each made at least marginally reliable contributions to this overall good fit: F_P, $t(12) = 2.21$, $p < .05$, and F_C, $t(12) = 1.36$, $p = .10$ (one-tailed). The root mean square error between the predicted and observed mean switching-time costs was approximately 70 ms. Two of the four predicted time costs differed by approximately 5% from the corresponding observed time costs, another one of them differed by approximately 10%, and none differed by as much as 20% (see Figure 11). The relative magnitudes of the predictor coefficients in Equation 1 suggest that prior-task familiarity may retard task switching more than current-task familiarity promotes it when visual pattern-classification tasks with low-complexity rules are involved.

In other respects, however, the mean switching-time costs obtained during Experiment 4 do not seem so consistent with the preceding pattern. For the high-rule-complexity condition, we found that switching from the size-and-shading classification task to the shape-and-numerosity classification task took reliably longer than switching in the opposite direction (mean difference = 303 ms, SE = 80 ms), $t(12) = 3.78$, $p < .005$. This is the reverse of what might be expected simply in terms of the task-familiarity hypothesis, given that decisions about the shape dimension are putatively more familiar than decisions about the size dimension, and decisions about the numerosity dimension are putatively more familiar than decisions about the shading dimension. Thus, our results suggest that some properties of rule activation for the visual pattern-classification tasks with low-complexity rules do not generalize to the visual pattern-classification tasks with high-complexity rules. Rather, in preparation for the latter tasks, the rule-activation stage may be more intricate than is characterized simply by task familiarity. Some possible sources of these additional intricacies are considered in the following discussion.

Discussion

By having participants perform various visual pattern-classification tasks in Experiment 4, we obtained more results that conform to predictions based on our stage model of executive control in task switching. Although the present procedure involved discrete RT trials, the mean time costs to switch between different tasks were reliably positive, and rule complexity affected them as in Experiment 1, where patterns of geometric objects were classified through continuous card sorting. Also as in Experiment 1, stimulus discriminability affected mean RTs reliably during both alternating-task and repetitive-task trial blocks, but it had virtually no effect on mean switching-time costs, which indicates that the discriminability effect was limited to a stimulus-identification

---

[11] Unfortunately, the RTs for the shading-discrimination and numerosity-discrimination tasks on repetitive-task blocks did not differ reliably in Experiment 4 (mean difference = 12 ms, SE = 24 ms), $t(12) = 0.5$, $p > .5$. Thus, unlike in Experiment 1, these dimensions could not be used to test our model's predictions about the effects of stimulus discriminability. This presumably happened because in Experiment 4, rather than manipulating the interior shading of the objects, we manipulated the shading of their backgrounds. Doing so helped make the levels of object shading equally discriminable regardless of object size, but it also eliminated the difference in degree of discriminability that occurred between these perceptual dimensions in Experiment 1.

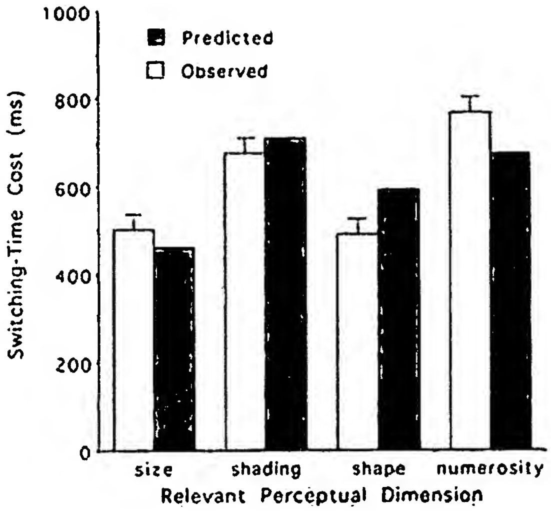

Figure 11. Predicted versus observed mean switching-time costs for the visual pattern-classification tasks that involved low-complexity rules in Experiment 4. Standard errors (lines extending above vertical time bars) are based on the interaction among trial-block type, relevant perceptual dimension, and participants. Scaled subjective familiarities of the prior and current tasks on alternating-task blocks provided the relevant predictor variables.

stage of task processes. Taken together, these results further substantiate the functional separability of executive control and task processes while reaffirming that executive control in task switching entails a rule-activation stage whose duration increases with rule complexity.

Experiment 4 likewise replicated and extended Experiment 3, yielding more evidence about asymmetries in task switching and the degree to which rule activation is affected by task familiarity. Under the low-rule-complexity condition of Experiment 4, switching from pattern-classification tasks that involved relatively familiar visual dimensions to ones that involved relatively unfamiliar visual dimensions was slower on average than switching in the opposite direction. This result resembles what occurred during Experiment 3, where the time costs of switching from relatively familiar to relatively unfamiliar types of arithmetic problems also tended to be greater than the time costs of switching in the opposite direction. Taken together, these results provide considerable support for the task-familiarity hypothesis, which claims that the rule-activation stage of executive control in task switching takes longer to disable the rules of prior familiar tasks or to enable the rules of current unfamiliar tasks.

More precisely, our multiple regression analysis (Equation 2) of the results from Experiment 4 revealed that, for the visual pattern-classification tasks whose rules had low complexity, mean switching-time costs were affected by both prior-task and current-task familiarity, but the detrimental effect of the former factor was greater than the beneficial effect of the latter factor. This quantitative difference offers deeper insights into how rule activation works under at least some conditions. Apparently, the operations of this executive-control stage differ from those through which the operating system of a modern digital computer accesses and initiates successive user programs from long-term (e.g., disk) storage (cf. Kieras et al., 2000). Rule activation involves more than just loading the current task's rules into procedural working memory while concomitantly overwriting the prior task's rules. Instead, the rules for the prior task may have to be suppressed during an inhibitory operation that is functionally distinct from the one that enables the rules for the current task (cf. Allport et al., 1994; Goschke, 2000; Mayr & Keele, 2000). Perhaps this inhibition is more time consuming for prior familiar tasks because when their rules are enabled in procedural working memory, they have a higher level of activation than do the rules of prior unfamiliar tasks. If so, then Experiment 4 not only helps to elaborate the task-familiarity hypothesis. It also raises further doubts about the general veracity of Allport et al.'s (1994) TSI hypothesis, which implies that because of TSI, unfamiliar-to-familiar switching should be more—not less—difficult than familiar-to-unfamiliar switching, contrary to our results from the low-rule-complexity condition.

Nevertheless, other subtleties may occur in the rule-activation stage for tasks that involve high-complexity rules. According to a straightforward extension of the task-familiarity hypothesis, switches from the bidimensional shape-and-numerosity classification task to the bidimensional size-and-shading classification task should have taken longer on average than switches in the opposite direction. This follows because participants had judged that visual decisions about the size dimension are less familiar than visual decisions about the shape dimension, and visual decisions about the shading dimension are less familiar than visual decisions about the numerosity dimension. However, during Experiment 4, we found that the magnitudes of the mean switching-time costs for the tasks in which stimuli had to be classified conjunctively with respect to these dimensions were reversed from what the task-familiarity hypothesis ordinarily would predict.

There are various conceivable explanations of this latter reversal. For example, consistent with the TSI hypothesis (Allport et al., 1994), perhaps sets of complex unfamiliar task rules require an extraordinary persistent boost in activation to be enabled, making it harder to perform a more familiar task subsequently after the rules of a related but less familiar task have received such a boost. This might explain why switching from the size-and-shading classification task to the shape-and-numerosity classification task took longer than switching in the opposite direction, given that visual decisions about the size and shading dimensions were judged to be respectively less familiar than visual decisions about the shape and numerosity dimensions.

Yet this explanation seems doubtful because our results from Experiment 3 tend to contradict it. In particular, we found that for arithmetic problems whose solutions require relatively complex (i.e., multiplication and division) rules, switching from the more familiar (i.e., multiplication) to the less familiar (i.e., division) problems took longer on average than switching in the opposite direction, contrary to the preceding account based on the TSI hypothesis. Thus, to us, it appears more likely that during Experiment 4, some other subtlety in the rule-activation stage caused the unexpected reversal of the mean switching-time costs for the bidimensional pattern-classification tasks.

These considerations lead to a second possibility. Thus far, our extension of the task-familiarity hypothesis to bidimensional pattern-classification tasks has assumed that they involve enabling two distinct sets of more or less familiar task rules, which would each be used for making a classification with respect to one relevant visual dimension, after which the respective unidimensional classifications would be combined. However, this assumption may not hold in all cases. Instead, some bidimensional pattern-classification tasks may be performed through their own specially formulated rules that take the configural features of stimuli directly into account. Such rules could be more compact and efficient than is afforded by a simple union of rule sets for unidimensional classifications. For example, perhaps size-and-shading classifications are made through special rules that take the amount of object shading into account without requiring a combination of separate decisions about object size and shading. On the other hand, unions of rule sets for making unidimensional classifications might still be used to perform bidimensional pattern-classification tasks whose relevant visual dimensions are not so perceptually integrable as size and shading. In particular, shape-and-numerosity classifications may be made through use of both a rule set for classifying object shapes and a rule set for classifying object numerosities that treat the shape and numerosity dimensions separately.

If so, this would explain why the mean time cost of switching from size-and-shading classifications to shape-and-numerosity classifications exceeds the mean time cost of switching in the opposite direction. Rule activation could take longer for shape-and-numerosity classifications because they involve enabling relatively many rules, in the form of two distinct rule sets, whereas size-and-shading classifications involve enabling fewer rules in the form of a single rule set. This time-cost difference may prevail even though, when made separately, visual decisions about the shape dimension and about the numerosity dimension are respectively more familiar than visual decisions about the size dimension and about the shading dimension. Thus, predictions regarding asymmetries in task switching based on task familiarity must accommodate supplementary context-dependent factors that arise when tasks whose rules are potentially complex have to be performed.

General Discussion

During the last decade of the 20th century, two contrasting types of theoretical proposals about the sources of time costs in task switching came to the fore (Monsell, 1996). According to one proposal, the time to complete a current task that immediately follows a different prior task is lengthened because persistent residual activation from one of the prior task's processes (e.g., response selection) interferes with some similar process for the current task. An influential representative of this type is Allport et al.'s (1994) TSI hypothesis. Also related to it is the hypothesis of contention scheduling in Norman and Shallice's (1986) ATA model. In both of these cases, the time costs of task switching are attributed to conflicts at a basic level of task processing without any assumption of control being exerted from a higher supervisory level. On the other hand, according to a second type of proposal, switching-time costs stem from the temporal overhead of executive control processes that reconfigure the settings of component information-processing mechanisms to be appropriate for performing successive tasks whose functional requirements are mutually exclusive. Representatives of this type include the supervisory attentional system of the ATA model, the endogenous and exogenous control processes hypothesized by Rogers and Monsell (1995), and the task-scheduling procedures in EPIC computational models of Kieras et al. (2000). In all of these latter cases, higher level supervision of task switching is assumed to be used for achieving successful performance, and switching-time costs are attributed to more than simply interference between the basic processes of successive tasks. Of course, the preceding alternative types of theoretical proposals could each be correct in part, but even if so, it is important to determine when, where, and to what extent functionally different sources contribute to the slowness of task switching under various conditions.

Given the latter objectives, the present four experiments provide informative data about the existence and nature of executive control in task switching. We found that in two different task domains, visual pattern classification and arithmetic problem solving, reliable mean switching-time costs occurred, and their magnitudes increased with the complexity of the rules needed for performing the tasks between which participants had to switch. Furthermore, there were systematic asymmetries in task switching; time costs tended to be greater on average when participants switched from a task that was relatively familiar or switched to a task that was relatively unfamiliar. Our overall findings suggest that task switching may often be mediated by a rule-activation stage of executive control through which the rules for prior tasks are disabled and the rules for current tasks are enabled in distinct operations.

Some of our additional findings suggest that such executive control may be functionally independent of accompanying basic task processes. For example, the discriminability of visual stimuli in the pattern-classification tasks affected mean RTs reliably during both alternating-task and repetitive-task blocks, but there were no reliable interactions between the discriminability and block-type effects. Instead, mean switching-time costs remained about the same regardless of stimulus discriminability, as if rule activation is separate from a specific task process in which the discriminability effect takes place selectively.

There likewise appear to be other executive control processes that are functionally separate from rule activation as well as basic task processes. In particular, we found that explicit task cues (arithmetic operation signs) reduced the mean time cost for switching between different types of arithmetic problems (i.e., addition and subtraction or multiplication and division). However, this task-cuing effect was approximately additive with the rule-complexity effect on mean switching-time costs, as would be expected if a distinct goal-shifting stage of executive control accompanies and complements the rule-activation stage.

Of course, the present results do not eliminate the possibility that, under some conditions, TSI or other such passive residual between-tasks interference effects contribute significantly to the time costs of task switching. Nevertheless, it is difficult for us to see how the TSI hypothesis alone could easily explain all of our findings. Thus, we construe our findings instead as revealing the overarching importance of executive control in task switching for at least the task domains studied here, and perhaps others as well.

More Evidence of Executive Control Processes in Task Switching

Some more evidence that executive control processes contribute to task switching and are separable from basic task processes has been reported by Monsell, Azuma, Eimer, Le Pelley, and Strafford (1998). Using Rogers and Monsell's (1995) alternating-runs paradigm, they had participants perform two tasks: vocally naming printed digits and classifying digits as odd or even by pressing left or right finger keys. On each trial of the odd–even classification task, a participant's RT and lateralized readiness potential (LRP) were recorded. The LRP is an event-related brain potential that manifests preparatory motor processes and response competition before overt manual movements (Coles, 1989; Osman, Bashore, Coles, Donchin, & Meyer, 1992). Given their functional significance, records of LRPs may reveal relationships among rule activation, response selection, and motor preparation on trials during which task switching occurs.

Exploiting this prospect, Monsell et al. (1998) compared RTs and latencies of LRP onsets for trials that involved task alternations and task repetitions. The alternating-task trials yielded longer mean RTs. There was a concomitant lengthening in the mean latency of stimulus-locked LRP onsets that was approximately equal to the mean switching-time cost calculated from the RT data, as if an extra control process had occurred sometime before motor preparation began on the alternating-task trials. However, response-locked LRPs on alternating-task and repetitive-task trials had virtually identical waveforms, unlike in other contexts where significant response competition has been found (e.g., Coles, 1989; Osman et al., 1992). The LRP waveforms obtained by Monsell et al. yielded no evidence of proactive interference with response selection on alternating-task trials, contrary to the TSI hypothesis of Allport et al. (1994). Instead, consistent with our stage model of executive control in task switching (see Figure 1), Monsell et al.'s results suggest that switching-time costs stemmed from a functionally and temporally separate rule-activation stage inserted between task processes of stimulus identification and response selection on alternating-task trials.

Other studies by Monsell et al. (2000) have further explored the conditions under which different types of asymmetry occur in task switching. Again contrary to the TSI hypothesis and results of Allport et al. (1994), Monsell et al. (2000) found a number of task combinations for which time costs were lower when participants switched from a relatively nondominant task (i.e., one with weak or incompatible S-R associations) to a relatively dominant task (i.e., one with strong or compatible S-R associations) instead of switching in the opposite direction. This natural asymmetry, favoring switches to the dominant rather than nondominant task, is analogous to our observations regarding the effects of prior-task and current-task familiarity on mean switching-time costs in Experiments 3 and 4. It is also analogous to what some other investigators have observed (e.g., Mayr & Kliegl, 2000).

Only in relatively rare cases did Monsell et al. (2000) find the converse paradoxical asymmetry that the TSI hypothesis predicts, in which mean time costs are greater for nondominant to dominant task switching than for dominant to nondominant task switching. Interestingly, these cases involved extremely large differences between the strengths of the S-R mappings for the dominant and nondominant tasks, and performance of the dominant tasks was considerably more automatized than performance of the nondominant tasks. Yet despite the special status of the latter circumstances, our stage model of executive control in task switching may be generalized to explain results from them as well.

Explanation of Paradoxical Asymmetries in Task Switching

To illustrate how our model can explain paradoxical asymmetries in task switching, we again consider Allport et al.'s (1994, Experiment 5) study. As mentioned before, they found a significant time cost for switching from the nondominant standard Stroop (ink-color naming) task to the dominant reverse Stroop (color-word reading) task, but there was essentially no time cost for switching from the reverse Stroop to the standard Stroop task. These findings suggest that, on alternating-task blocks, goal shifting and rule activation perhaps contributed to RTs for the reverse Stroop task, whereas RTs for the standard Stroop task did not include such contributions.

What underlies this paradoxical asymmetry? Following previous speculations, one possibility is that the production rules for some extremely dominant tasks are permanently enabled in procedural long-term memory, which—counterintuitively—can lead the durations of goal shifting and rule activation to be obscured during switches to other nondominant tasks (cf. Monsell et al., 2000). Specifically, consistent with some theories about performance of the standard Stroop task (e.g., Kornblum et al., 1990; MacLeod, 1991; Posner & Snyder, 1975; Schweickert, 1978), let us augment our model with the following assumptions. First, because of extensive prior practice, procedural long-term memory contains permanently enabled (active) rules that select phonological response codes for vocal reading of printed words. These word-reading rules are applied automatically and obligatorily whenever a printed word has been perceived; they do not have to be enabled by an intentional rule-activation stage. Second, there is another ordinarily disabled (inactive) set of rules that select phonological response codes for vocal naming of ink colors. Application of the ink-color naming rules is optional; an intentional rule-activation stage must enable them before they can be applied. Third, performance of the standard Stroop task relies on the ink-color naming rules and a set of supplementary color-word editing rules. Fourth, during switches from the reverse Stroop to the standard Stroop task, the rule-activation stage enables the ink-color naming and color-word editing rules, which are applied as soon as possible thereafter. Fifth, the color-word editing rules discard irrelevant phonological response codes that have been selected through color-word reading, permitting correct vocalization of ink-color names that are selected through the ink-color naming rules. Sixth, RTs for the standard Stroop task stem from a competitive race that involves ink-color identification and response selection versus color-word reading and editing. On each trial, the durations of these two processing sequences determine the observed RT; overt responses occur only after the last competitor finishes the race. Finally, during alternating-task trial blocks, the goal-shifting and rule-activation stages of executive control also participate in this race, but they may not be on the "critical path" to an overt response.

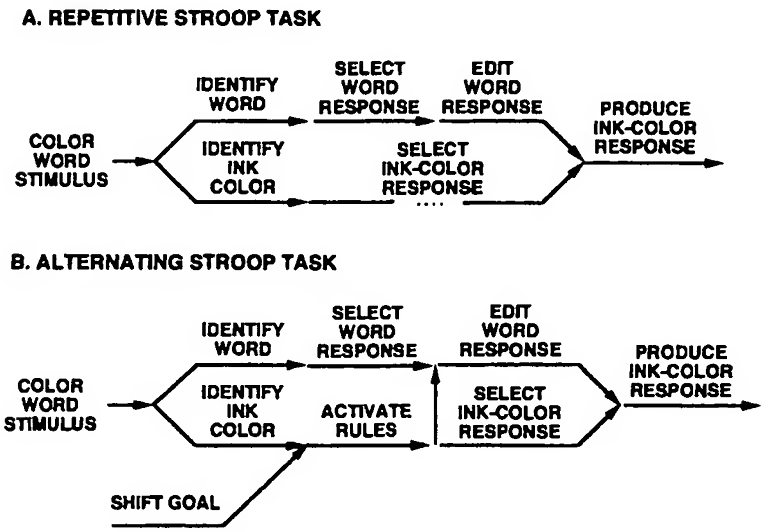

Figure 12. Order-of-processing diagrams with temporal relationships among executive control and task processes for the standard Stroop task. Processes that proceed along different pathways occur in parallel. Processes that proceed along the same pathway occur in sequence. Each process begins only after all processes that precede it in sequence have ended. A: Sequential and parallel processes on repetitive-task blocks. B: Sequential and parallel processes on alternating-task blocks in which the standard Stroop task has been preceded by the reverse Stroop task.

Null switching-time cost for standard Stroop task. Now augmented by these assumptions, our model provides a quantitative account of the null switching-time cost that Allport et al. (1994, Experiment 5) found for the standard Stroop task. To see how, we must examine the standard Stroop RTs in detail during repetitive-task and alternating-task blocks, respectively. For this purpose, some order-of-processing diagrams of the various processing sequences that presumably mediate these RTs appear in Figure 12 (cf. Fisher & Goldstein, 1983; Schweickert, 1978).

On each trial of a repetitive-task block with the standard Stroop task (Figure 12A), the theoretical RT according to our model and present assumptions is as follows:

$$\mathrm{RT} = \max(t_{c1} + t_{c2}, t_{w1} + t_{w2} + t_e) + t_{c3}. \quad (3)$$

Here t_{c1} is the time taken to perceptually identify an ink color; t_{c2} is the time taken to select a relevant ink-color response code; t_{w1} is the time taken to perceptually identify a printed color word; t_{w2} is the time taken to select an irrelevant color-word response code; and t_e is the time taken by the response-editing rules to discard the irrelevant code. Color-word response selection and editing presumably contribute to Equation 3 for reasons mentioned before (i.e., responses to color words are assumed to be selected obligatorily and must be edited because they could disrupt ink-color naming). Combining these and other components in this equation, $\max(t_{c1} + t_{c2}, t_{w1} + t_{w2} + t_e)$ is the total time taken to complete the race that involves ink-color naming versus color-word reading and editing, and t_{c3} is the time taken to produce an overt vocal response based on the relevant ink-color response code after the race has finished.[12] Of course, consistent with standard Stroop interference (Allport et al., 1994; MacLeod, 1991), this combination yields longer RTs than would occur in a control condition that requires naming the ink colors of meaningless alphanumeric character strings (e.g., XXXX), for which the expected RT is simply $t_{c1} + t_{c2} + t_{c3}$.

Be this as it may, the situation is even more complicated when switches must occur from the reverse Stroop to the standard Stroop task (see Figure 12B). During each of these switches, automatic obligatory color-word reading would proceed as before (cf. Figure 12A). Meanwhile, concomitant with it, there would be a race between goal shifting and ink-color identification whose joint completion determines when activation of the color-word editing and ink-color response-selection rules begins. Also, following the rule-activation stage, there would be a race between the editing of an irrelevant color-word response code and the selection of a relevant ink-color response code. As a result, because of the partial temporal overlap among these various processing sequences, goal shifting and rule activation may fall off the critical path whose total duration determines overall RT (cf. Schweickert, 1978).

[12] The theoretical RT for the baseline task of repetitively naming ink-color patches is simply $\mathrm{RT} = t_{c1} + t_{c2} + t_{c3}$. Given Equation 3, the standard Stroop task can manifest interference relative to the baseline task because $t_{w1} + t_{w2} + t_e$ may exceed $t_{c1} + t_{c2}$, as Allport et al. (1994, Experiment 5) and others (e.g., MacLeod, 1991) have found. Editing of the irrelevant color-word response code, which contributes t_e to Equation 3, must be included because color-word reading occurs automatically and is typically faster than ink-color naming (i.e., $t_{w1} + t_{w2} < t_{c1} + t_{c2}$).

More precisely, on each trial of an alternating-task block with the standard Stroop task and a zero RSI, the theoretical RT according to our model and present assumptions is as follows:

$$\mathrm{RT} = \max[t_{sc} + t_{c2}, \max(t_{sc}, t_{w1} + t_{w2}) + t_e] + t_{c3}. \quad (4)$$

Here the terms t_{c2}, t_{c3}, t_{w1}, t_{w2}, and t_e are the same as in Equation 3. However, Equation 4 also has some new terms, and it contains contributions from a number of processes that are partly serial and partly parallel (see Figure 12B).

For example, one major component of Equation 4 is the sum $t_{sc} + t_{c2}$. As part of it, t_{sc} is the time taken from the onset of the stimulus (i.e., color word printed in colored ink) until the rule-activation stage for ink-color naming and color-word response editing has finished. By definition, $t_{sc} = \max(t_g, t_{c1}) + t_{ac}$, where t_g is the duration of the goal-shifting stage, t_{c1} is the duration of the perceptual ink-color identification stage (cf. Equation 3), and t_{ac} is the duration of the rule-activation stage for ink-color naming. The term $\max(t_g, t_{c1})$ contributes to t_{sc} because rule activation is assumed to start only after both goal shifting and ink-color identification have finished. Furthermore, t_{c2} is the time taken to select a relevant ink-color response code (cf. Equation 3). Thus, the sum $t_{sc} + t_{c2}$ in Equation 4 is the duration of the processing sequence that yields a relevant ink-color response code through goal shifting, color identification, rule activation, and ink-color response selection.

A second major component of Equation 4 is $\max(t_{sc}, t_{w1} + t_{w2}) + t_e$, which represents the duration of another temporally overlapped processing sequence. It involves automatic color-word identification and color-word response selection followed by editing of the irrelevant color-word response code after the rule-activation stage has enabled the editing rules (see Figure 12B). Consequently, the overall expression $\max[t_{sc} + t_{c2}, \max(t_{sc}, t_{w1} + t_{w2}) + t_e]$ on the right side of Equation 4 embodies races among executive control, irrelevant automatic word reading, and relevant controlled color naming. In turn, the RT for the standard Stroop task during alternating-task blocks depends on which time component, $t_{sc} + t_{c2}$ or $\max(t_{sc}, t_{w1} + t_{w2}) + t_e$, is greater.

Given this dependence, we see that the durations of goal shifting and rule activation may not contribute to switching-time costs for the standard Stroop task. In particular, suppose that the following inequalities hold: $t_g \le t_{c1}$; $t_{ac} \le t_{w1} + t_{w2} - t_{c1}$; $t_{w1} + t_{w2} < t_{c1} + t_{c2}$; and $t_e \ge t_{ac} + t_{c1} + t_{c2} - t_{w1} - t_{w2}$. Then the right side of Equation 4 reduces to $\max(t_{c1} + t_{c2}, t_{w1} + t_{w2} + t_e) + t_{c3}$, and there would be no apparent time cost of task switching. Instead, the theoretical RTs for the standard Stroop task during alternating-task blocks would equal those during repetitive-task blocks (cf. Equation 3), consistent with Allport et al. (1994, Experiment 5). A null switching-time cost can occur even though goal shifting and rule activation mediate switches to this task. Allport et al.'s (1994, Experiment 5) results do not prove that executive control is absent under these conditions, because contributions of crucial control stages can be obscured by the durations of other concomitant processes.

Positive switching-time cost for reverse Stroop task. Through a similar rationale, the present model also provides an account of why Allport et al. (1994, Experiment 5) found a positive switching-time cost for the reverse Stroop (i.e., color-word reading) task. This account holds even though the reverse Stroop task is presumably much more dominant than the standard Stroop (i.e., ink-color naming) task. To see how, we must examine the reverse Stroop RTs in detail during repetitive-task and alternating-task blocks, respectively. For this purpose, some order-of-processing diagrams of the various steps that presumably mediate these RTs appear in Figure 13.

On each trial of a repetitive-task block with the reverse Stroop task (see Figure 13A), the theoretical RT according to our model and present assumptions is as follows:

$$\mathrm{RT} = t_{w1} + t_{w2} + t_{w3}, \quad (5)$$

where the terms on the right side (i.e., times for perceptual color-word identification, color-word response selection, and overt vocalization) are the same as before (cf. Equations 3 and 4).[13] No contributions from irrelevant ink-color naming appear in Equation 5 because the rules for ink-color naming are presumably disabled during blocks of trials that require only color-word reading. Thus, consistent with Allport et al. (1994, Experiment 5), our model implies that the reverse Stroop task may manifest no interference in repetitive-task blocks relative to an appropriate control condition (i.e., reading color words printed in black ink).

Again, however, the situation is more complicated when switches must be made from the standard Stroop to reverse Stroop task (see Figure 13B). During each of these switches, automatic color-word reading would proceed as before (cf. Figure 13A). Meanwhile, concomitant with it, there would be a race between goal shifting and perceptual color-word identification, whose joint completion determines when rule activation begins for the reverse Stroop task. Rule activation is required in this case because the rules used previously to select ink-color response codes must be disabled, and other rules used to release the next relevant color-word response code for overt vocalization must be enabled. After completion of this rule-activation stage and selection of a color-word response, the response code would be released by the currently enabled rules and then vocalized overtly.

More precisely, on each trial of an alternating-task block with the reverse Stroop task and a zero RSI, the theoretical RT according to our model and present assumptions is as follows:

$$\mathrm{RT} = \max(t_{sw}, t_{w1} + t_{w2}) + t_{w3}. \quad (6)$$

In this equation, t_{sw} is the time taken by a processing sequence that involves goal shifting, perceptual color-word identification, and activation of the rules for releasing selected color-word responses (cf. Figure 13B). By definition, $t_{sw} = \max(t_g, t_{w1}) + t_{aw}$, where t_g is the duration of the goal-shifting stage, t_{w1} is the duration of the color-word identification stage, and t_{aw} is the duration of the rule-activation stage. The term $\max(t_g, t_{w1})$ contributes to t_{sw} because rule activation is assumed to start only after both goal shifting and color-word identification have finished. Paralleling these operations, $t_{w1} + t_{w2}$ is the duration of another processing sequence that involves automatic color-word identification and color-word response selection. Consequently, the overall expression $\max(t_{sw}, t_{w1} + t_{w2})$ on the right side of Equation 6 embodies a race between executive control and the automatic phase of

[13] In particular, we assume here that $t_{w3} = t_{c3}$.

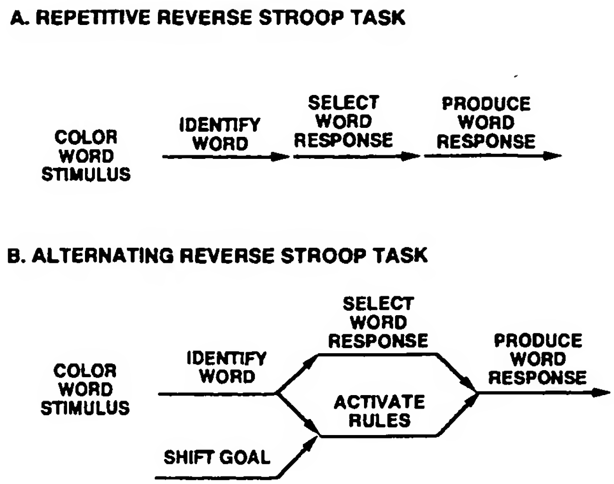

Figure 13. Order-of-processing diagrams with temporal relationships among executive control and task processes for the reverse Stroop task. Processes that proceed along different pathways occur in parallel. Processes that proceed along the same pathway occur in sequence. Each process begins only after all processes that precede it in sequence have ended. A: Sequential and parallel processes on repetitive-task blocks. B: Sequential and parallel processes on alternating-task blocks in which the reverse Stroop task has been preceded by the standard Stroop task.

color-word reading for the reverse Stroop task during alternating-task blocks. In turn, the RT for the reverse Stroop task during alternating-task blocks depends on which time component, t_{rw} or $t_{w1} + t_{w2}$, is greater.

Given this dependence, we see again that contributions of goal shifting and rule activation to theoretical RTs may be obscured by the durations of other concomitant processes. Nevertheless, if $t_g > t_{w1}$ and $t_{aw} > t_{w2}$, or if only $t_{aw} > t_{w2}$, then the right side of Equation 6 would exceed the right side of Equation 5. As a result, there could be a significant positive switching-time cost for the reverse Stroop task, consistent with Allport et al. (1994, Experiment 5). Indeed, according to our model, both this latter result and a null switching-time cost for the standard Stroop task can occur concomitantly, yielding the paradoxical asymmetry in task switching that originally inspired Allport et al.'s TSI hypothesis. Furthermore, our model can likewise explain additional paradoxical asymmetries reported by Monsell et al. (2000), because they too involved cases in which one task was considerably more dominant than another.

Directions for Future Research

The preceding considerations point toward several promising directions for future research. To obtain a more complete description of the control processes that mediate task-set reconfiguration, further studies should be conducted with a large variety of tasks in the successive-tasks procedure. As part of these studies, multiple relevant independent variables (e.g., task familiarity, rule complexity, operation cuing, and RSI) should be manipulated orthogonally, and RTs should be analyzed to discover patterns of additive and interactive factor effects on switching-time costs as well as other related dependent variables. Following from the methodology and results of our present experiments, this detailed analytical approach promises to yield more insights about stages of executive control in task switching.

Along the way, concerted efforts should be made to investigate how practice and systematic training protocols change the control processes whereby task switching is accomplished. We and others (e.g., Allport et al., 1994; Mayr & Kliegl, 2000; Monsell et al., 2000) have found that the durations of these processes are affected by task dominance, familiarity, and other factors related to the degree of automaticity in task performance. Through exploring such effects more thoroughly, additional informative tests of the TSI hypothesis, the task-familiarity hypothesis, and alternative theoretical proposals about adaptive executive control will be possible (cf. Lauber, 1995; Schumacher et al., 1999, 2001). On the basis of results from these tests, precise veridical computational models of executive control and task switching can ultimately be formulated (Kieras et al., 2000).

References

Allport, A. (1980). Patterns and actions: Cognitive mechanisms are content-specific. In G. L. Claxton (Ed.), *Cognitive psychology: New directions* (pp. 26–64). London: Routledge & Kegan Paul.

Allport, A., Styles, E., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 421–452). Cambridge, MA: MIT Press.

Allport, A., & Wylie, G. (2000). Task switching, stimulus-response bindings, and negative priming. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 35–70). Cambridge, MA: MIT Press.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Anderson, S., Damasio, H., Jones, R., & Tranel, D. (1991). Wisconsin Card Sorting Test performance as a measure of frontal-lobe damage. *Journal of Clinical and Experimental Neuropsychology, 13,* 909–922.

Baddeley, A. D. (1986). *Working memory.* Oxford, England: Oxford University Press.

Bertelson, P. (1966). Central intermittency 20 years later. *Quarterly Journal of Experimental Psychology, 18,* 153–164.

Biederman, I. (1973). Mental set and mental arithmetic. *Memory & Cognition, 1,* 383–386.

Botwinick, J., Brinley, J. F., & Robbin, J. S. (1958). Task alternation time in relation to problem difficulty and age. *Journal of Gerontology, 13,* 414–417.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction.* Hillsdale, NJ: Erlbaum.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97,* 332–361.

Coles, M. G. H. (1989). Modern mind-brain reading: Psychophysiology, physiology, and cognition. *Psychophysiology, 26,* 251–269.

Coombs, C. H. (1964). *A theory of data.* New York: Wiley.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Erlbaum.

Dark, V. (1990). Switching between memory and perception: Moving attention or moving memory retrieval. *Memory & Cognition, 18,* 119–127.

Donders, F. C. (1969). On the speed of mental processes (W. G. Koster,

Trans.). In W. G. Koster (Ed.), *Attention and performance II* (pp. 412–431). Amsterdam: North-Holland. (Original work published 1868)
Drewe, E. A. (1974). The effect of type and area of brain lesion on Wisconsin Card Sorting Test performance. *Cortex, 10,* 159–170.
Duncan, J. (1986). Disorganization of behaviour after frontal-lobe damage. *Cognitive Neuropsychology, 3,* 271–290.
Fisher, D. L., & Goldstein, W. M. (1983). Stochastic PERT networks as models of cognition: Derivation of mean, variance, and distribution of reaction time using order-of-processing (OP) diagrams. *Journal of Mathematical Psychology, 27,* 121–151.
Garcia-Ogueta, M. I. (1993). Internal attentional switching: Effects of predictability, complexity and practice. *Acta Psychologica, 83,* 13–32.
Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1993). A neural system for error-detection and compensation. *Psychological Science, 4,* 385–390.
Gopher, D., Armony, L., & Greenshpan, Y. (2000). Switching tasks and attention policies. *Journal of Experimental Psychology: General, 129,* 308–339.
Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task-set switching. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 331–355). Cambridge, MA: MIT Press.
Grafman, J., Jones, B., & Salazar, A. (1990). Wisconsin Card Sorting Test performance based on location and size of neuroanatomical lesion in Vietnam veterans with penetrating head injury. *Perceptual and Motor Skills, 71,* 1120–1122.
Grant, A. D., & Berg, E. A. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology, 38,* 404–411.
Hawkins, H. L., Rodriguez, E., & Reicher, G. M. (1979). *Is time-sharing a general ability?* (Tech. Rep. No. 3). Eugene: University of Oregon.
Hays, W. L. (1963). *Statistics for psychologists.* New York: Holt, Rinehart & Winston.
Heaton, R. (1981). *A manual for the Wisconsin Card Sorting Test.* Odessa, FL: Psychological Assessment Resources.
Hunt, E., & Lansman, M. (1986). A unified model of attention and problem solving. *Psychological Review, 93,* 446–461.
James, W. (1890). *The principles of psychology.* New York: Holt.
Jersild, A. (1927). Mental set and shift. *Archives of Psychology* (Whole No. 89).
Kantowitz, B. H. (1974). Double stimulation. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 83–131). Hillsdale, NJ: Erlbaum.
Karlin, L., & Kestenbaum, R. (1968). Effects of number of alternatives on the psychological refractory period. *Quarterly Journal of Experimental Psychology, 20,* 67–178.
Keele, S. W., & Hawkins, H. L. (1982). Explorations of individual differences relevant to high level skill. *Journal of Motor Behavior, 14,* 3–23.
Keele, S. W., & Rafal, R. D. (2000). Deficits of attentional set in patients with left prefrontal cortex lesions. In S. Monsell & J. Driver (Eds.), *Attention and Performance XVIII: Control of cognitive processes* (pp. 627–651). Cambridge, MA: MIT Press.
Kieras, D. E., Meyer, D. E., Ballas, J. A., & Lauber, E. J. (2000). Modern computational perspectives on executive mental control: Where to from here? In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 681–712). Cambridge, MA: MIT Press.
Kimberg, D. Y., & Farah, M. J. (1993). A unified account of cognitive impairments following frontal lobe damage: The role of working memory in complex organized behavior. *Journal of Experimental Psychology: General, 122,* 411–428.
Kornblum, S., Hasbroucq, T., & Osman, A. (1990). Dimensional overlap: Cognitive basis of stimulus-response compatibility—A model and taxonomy. *Psychological Review, 97,* 253–270.
LaBerge, D., Petersen, R. J., & Norden, M. J. (1977). Exploring the limits of cueing. In S. Dornik (Ed.), *Attention and performance VI* (pp. 285–306). Hillsdale, NJ: Erlbaum.
Lauber, E. J. (1995). *Executive control of task switching operations.* Unpublished doctoral dissertation, University of Michigan, Ann Arbor.
Logan, G. (1985). Executive control of thought and action. *Acta Psychologica, 60,* 193–210.
Logan, G., & Zbrodoff, N. J. (1982). Constraints on strategy construction in a speeded discrimination task. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 502–520.
Los, S. A. (1996). On the origin of mixing costs: Exploring information processing in pure and mixed trial blocks. *Acta Psychologica, 94,* 144–188.
Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization.* New York: Oxford University Press.
MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109,* 163–203.
Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General, 129,* 4–26.
Mayr, U., & Kliegl, R. (2000). Task-set switching and long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1124–1140.
McCann, R. S., & Johnston, J. C. (1992). Locus of the single-channel bottleneck in dual-task performance. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 471–484.
McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review, 86,* 287–330.
Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1423–1442.
Meiselman, K. (1974). Time taken to switch attention between modalities as a function of mental effort. *Perceptual and Motor Skills, 39,* 1043–1046.
Meyer, D. E., Evans, J. E., Lauber, E. J., Gmeindl, L., Rubinstein, J., Junck, L., & Koeppe, R. A. (1998, April). *The role of dorsolateral prefrontal cortex for executive cognitive processes in task switching.* Poster presented at the meeting of the Cognitive Neuroscience Society, San Francisco, CA.
Meyer, D. E., Evans, J. E., Lauber, E. J., Rubinstein, J., Gmeindl, L., Junck, L., & Koeppe, R. A. (1997, March). *Activation of brain mechanisms for executive mental processes in cognitive task switching.* Poster presented at the meeting of the Cognitive Neuroscience Society, Boston, MA.
Meyer, D. E., & Kieras, D. E. (1997a). EPIC—A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104,* 3–65.
Meyer, D. E., & Kieras, D. E. (1997b). EPIC—A computational theory of executive cognitive processes and multiple-task performance: Part 2. Accounts of psychological refractory-period phenomena. *Psychological Review, 104,* 749–791.
Meyer, D. E., & Kieras, D. E. (1999). Précis to a practical unified theory of cognition and action: Some lessons from computational modeling of human multiple-task performance. In D. Gopher & A. Koriat (Eds.), *Attention and performance XVII: Cognitive regulation of performance* (pp. 17–88). Cambridge, MA: MIT Press.
Meyer, D. E., Kieras, D. E., Lauber, E., Schumacher, E., Glass, J., Zurbriggen, E., Gmeindl, L., & Apfelblat, D. (1995). Adaptive executive control: Flexible multiple-task performance without pervasive immutable response-selection bottlenecks. *Acta Psychologica, 90,* 163–190.

Meyer, D. E., Osman, A. M., Irwin, D. E., & Yantis, S. (1988). Modern mental chronometry. *Biological Psychology, 26,* 3–67.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1975). Loci of contextual effects on visual word-recognition. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and performance V* (pp. 98–135). London: Academic Press.

Miller, J. (1988). Discrete and continuous models of human information processing: Theoretical distinctions and empirical results. *Acta Psychologica, 67,* 191–257.

Milner, B. (1963). Effects of different brain lesions on card-sorting. *Archives of Neurology, 9,* 90–100.

Monsell, S. (1996). Control of mental processes. In V. Bruce (Ed.), *Unsolved mysteries of the mind: Tutorial essays in cognition* (pp. 93–148). Hove, England: Erlbaum, Taylor, & Francis.

Monsell, S., Azuma, R., Eimer, M., Le Pelley, M., & Strafford, S. (1998, July). *Does a prepared task switch require an extra (control) process between stimulus onset and response selection?* Poster presented at the 18th International Symposium on Attention and Performance, Windsor Great Park, United Kingdom.

Monsell, S., Yeung, N., & Azuma, R. (2000). Reconfiguration of task-set: Is it easier to switch to the weaker task? *Psychological Research, 63,* 250–264.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

Newell, A. (1973). Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing* (pp. 463–526). New York: Academic Press.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–18). New York: Plenum.

Osman, A., Bashore, T. R., Coles, M. G. H., Donchin, E., & Meyer, D. E. (1992). On the transmission of partial information: Inferences from movement-related brain potentials. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 217–232.

Pachella, R. G. (1974). The interpretation of reaction time in information processing. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 41–82). Hillsdale, NJ: Erlbaum.

Pashler, H. (1984). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 358–377.

Pashler, H. (1990). Do response modality effects support multiprocessor models of divided attention? *Journal of Experimental Psychology: Human Perception and Performance, 16,* 826–842.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116,* 220–244.

Pashler, H., & Johnston, J. C. (1989). Chronometric evidence of central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 41*(A), 19–45.

Posner, M. I., & Snyder, C. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition* (pp. 55–85). Hillsdale, NJ: Erlbaum.

Reason, J. T. (1990). *Human error.* Cambridge, England: Cambridge University Press.

Roberts, S., & Sternberg, S. (1993). The meaning of additive reaction-time effects: Tests of three alternatives. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 611–653). Cambridge, MA: MIT Press.

Rogers, R., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124,* 207–231.

Rubinstein, J. S. (1993). *Executive control of cognitive processes in task switching.* Unpublished doctoral dissertation, University of Michigan, Ann Arbor.

Rubinstein, J. S., Evans, J. E., & Meyer, D. E. (1994, March). *Task switching in patients with prefrontal cortex damage.* Poster presented at the meeting of the Cognitive Neuroscience Society, San Francisco, CA.

Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (1994, November). *Executive control of cognitive processes in task switching.* Poster presented at the meeting of the Psychonomic Society, St. Louis, MO.

Rumelhart, D. E., & Norman, D. A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science, 6,* 1–36.

Sanders, A. F. (1980). Stage analysis of reaction processes. In G. E. Stelmach & J. Requin (Eds.), *Tutorials in motor behavior* (pp. 331–354). Amsterdam: North-Holland.

Schumacher, E. H., Lauber, E. J., Glass, J. M. B., Zurbriggen, E. L., Gmeindl, L., Kieras, D. E., & Meyer, D. E. (1999). Concurrent response-selection processes in dual-task performance: Evidence for adaptive executive control of task scheduling. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 791–814.

Schumacher, E. H., Seymour, T. L., Glass, J. M., Fencsik, D., Lauber, E. J., Kieras, D. E., & Meyer, D. E. (2001). Virtually perfect time-sharing in dual-task performance: Uncorking the central cognitive bottleneck. *Psychological Science, 12,* 101–108.

Schweickert, R. (1978). A critical path generalization of the additive factor method: Analysis of a Stroop task. *Journal of Mathematical Psychology, 18,* 105–139.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society London, B298,* 199–209.

Shallice, T. (1988). *From neuropsychology to mental structure.* Cambridge, England: Cambridge University Press.

Shallice, T. (1994). Multiple levels of control processes. In C. Umilta & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 395–420). Cambridge, MA: MIT Press.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127–190.

Siegel, S. (1956). *Nonparametric statistics for the behavioral sciences.* New York: McGraw-Hill.

Smith, M. C. (1967). Theories of the psychological refractory period. *Psychological Bulletin, 67,* 202–213.

Spector, A., & Biederman, I. (1976). Mental set and mental shift revisited. *American Journal of Psychology, 89,* 669–679.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica, 30,* 276–315.

Sternberg, S. (1998). Discovering mental processing stages: The method of additive factors. In D. Scarborough & S. Sternberg (Eds.), *An invitation to cognitive science: Vol. 4. Methods, models, and conceptual issues* (pp. 703–863). Cambridge, MA: MIT Press.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18,* 643–662.

Stuss, D. T., & Benson, D. F. (1986). *The frontal lobes.* New York: Raven Press.

Sudevan, P., & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human Perception and Performance, 13,* 89–103.

Townsend, J. T. (1986). *Toward a dynamic mathematical theory of mental workload in POPCORN.* West Lafayette, IN: Behavioral Institute for Technology and Science.

Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes.* Cambridge, England: Cambridge University Press.

Weber, R., Burt, D., & Noll, N. (1986). Attention switching between perception and memory. *Memory & Cognition, 14,* 238–245.

Welford, A. T. (1952). The "psychological refractory period" and the timing of high speed performance—A review and a theory. *British Journal of Psychology, 43,* 2–19.

Welford, A. T. (1959). Evidence of a single-channel decision mechanism limiting performance in a serial reaction task. *Quarterly Journal of Experimental Psychology, 2,* 193–210.

Welford, A. T. (1967). Single channel operation in the brain. *Acta Psychologica, 27,* 5–22.

Received December 18, 1996
Revision received January 5, 2001
Accepted January 13, 2001 ■